

EXHIBIT NN

**UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

VIRTAMOVE, CORP.,

Plaintiff,

v.

INTERNATIONAL BUSINESS
MACHINES CORP.,

Defendant.

Case No. 2:24-cv-00064-JRG-RSP

**PLAINTIFF VIRTAMOVE, CORP.’S CORRECTED PRELIMINARY DISCLOSURE
OF ASSERTED CLAIMS AND INFRINGEMENT CONTENTIONS**

I. Patent Rule 3-1: Disclosure of Asserted Claims and Infringement Contentions

Pursuant to Patent Rule 3-1, Plaintiff VirtaMove, Corp. submits the following Preliminary Disclosure of Asserted Claims and Infringement Contentions. This disclosure is based on the information available to VirtaMove as of the date of this disclosure, and VirtaMove reserves the right to amend this disclosure to the full extent permitted, consistent with the Court’s Rules and Orders.

A. Patent Rule 3-1(a): Asserted Claims

VirtaMove asserts that Defendant International Business Machines Corp. (“Defendant” or “IBM”) infringes the following claims (collectively, “Asserted Claims”):

- (1) U.S. Patent No. 7,519,814 (“the ’814 patent”), claims 1, 2, 6, 9, and 10; and
- (2) U.S. Patent No. 7,784,058 (“the ’058 patent”), claims 1–4 and 18.

This Corrected Preliminary Disclosure of Asserted Claims and Infringement Contentions correctly reflects, consistent with the Complaint (IBM Dkt. 1) and the Amended Complaints (Consolidated Dkts. 37, 47), that the only independent claims that are asserted in this case are

independent claim 1 of the '814 patent and independent claim 1 of the '058 patent. Independent Claim 31 of the '814 patent is not, was not, and will not be asserted in this case. Otherwise, this Corrected Preliminary Disclosure of Asserted Claims and Infringement Contentions is identical to the previously served Preliminary Disclosure of Asserted Claims and Infringement Contentions.

B. Patent Rule 3-1(b): Accused Instrumentalities of which VirtaMove is aware

VirtaMove asserts that the Asserted Claims are infringed by the various instrumentalities used, made, sold, offered for sale, or imported into the United States by Defendant, including certain (a) IBM products and services using secure containerized applications, including without limitation IBM's Cloud Kubernetes Service (IKS), IBM Cloud Private (ICP), and IBM Hybrid Cloud mesh, and all versions and variations thereof since the issuance of the '814 patent; and (b) IBM products and services using user mode critical system elements as shared libraries, including without limitation IBM Cloud Kubernetes Service (IKS), IBM Cloud Private (ICP), and IBM Hybrid Cloud mesh, and all versions and variations thereof since the issuance of the '058 patent ("Accused Instrumentalities"). Defendant's Accused Instrumentalities of which VirtaMove is presently aware are described in more detail in the accompanying preliminary infringement contention charts.

VirtaMove reserves the right to accuse additional products from Defendant to the extent VirtaMove becomes aware of additional products during the discovery process. Unless otherwise stated, VirtaMove's assertions of infringement apply to all variations, versions, and applications of each of the Accused Instrumentalities, on information and belief, that different variations, versions, and applications of each of the Accused Instrumentalities are substantially the same for purposes of infringement of the Asserted Claims.

C. Patent Rule 3-1(c): Claim Charts

VirtaMove's analysis of Defendant's products is based upon limited information that is publicly available, and based on VirtaMove's own investigation prior to any discovery in these actions. Specifically, VirtaMove's analysis is based on certain limited resources that evidence certain products made, sold, used, or imported into the United States by Defendant.

VirtaMove reserves the right to amend or supplement these disclosures for any of the following reasons:

- (1) Defendant and/or third parties provide evidence relating to the Accused Instrumentalities;
- (2) VirtaMove's position on infringement of specific claims may depend on the claim constructions adopted by the Court, which has not yet occurred; and
- (3) VirtaMove's investigation and analysis of Defendant's Accused Instrumentalities is based upon public information and VirtaMove's own investigations. VirtaMove reserves the right to amend these contentions based upon discovery of non-public information that VirtaMove anticipates receiving during discovery.

Attached, and incorporated herein in their entirety, are charts identifying where each element of the Asserted Claims are found in the Accused Instrumentalities.

Unless otherwise indicated, the information provided that corresponds to each claim element is considered to indicate that each claim element is found within each of the different variations, versions, and applications of each of the respective Accused Instrumentalities described above.

D. Patent Rule 3-1(d): Literal Infringement / Doctrine of Equivalents

With respect to the patents at issue, each element of each Asserted Claim is considered to be literally present. VirtaMove also contends that each Asserted Claim is infringed or has been infringed under the doctrine of equivalents in Defendant's Accused Instrumentalities. VirtaMove

also contends that Defendant both directly and indirectly infringes the Asserted Claims. For example, the Accused Instrumentalities are provided by the Defendant to customers, who are actively encouraged and instructed (for example, through Defendant's online instructions on its website and instructions, manual, or user guides that are provided with the Accused Instrumentalities) by Defendant to use the Accused Instrumentalities in ways that directly infringe the Asserted Claims. Defendant therefore specifically intends for and induces its customers to infringe the Asserted Claims under Section 271(b) through the customers' normal and customary use of the Accused Instrumentalities. In addition, Defendant is contributorily infringing the Asserted Claims under Section 271(c) and/or Section 271(f) by selling, offering for sale, or importing the Accused Instrumentalities into the United States, which constitute a material part of the inventions claimed in the Asserted Claims, are especially made or adapted to infringe the Asserted Claims, and are otherwise not staple articles or commodities of commerce suitable for non-infringing use.

E. Patent Rule 3-1(e): Priority Dates

The Asserted Claims of the '814 patent are entitled to a priority date at least as early as September 15, 2003, the filing date of provisional application No. 60/502,619.

The Asserted Claims of the '058 patent are entitled to a priority date at least as early as September 22, 2003, the filing date of provisional application No. 60/504,213.

A diligent search continues for additional responsive information and VirtaMove reserves the right to supplement this response.

F. Patent Rule 3-1(f): Identification of Instrumentalities Practicing the Claimed Invention

At this time, VirtaMove does not identify any of its instrumentalities as practicing the Asserted Claims. A diligent search continues for additional responsive information and VirtaMove reserves the right to supplement this response.

II. Patent Rule 3-2: Document Production Accompanying Disclosure

Pursuant to Patent Rule 3-2, VirtaMove submitted the following Document Production Accompanying Disclosure, along with an identification of the categories to which each of the documents corresponds.

F. Patent Rule 3-2(a) documents:

VirtaMove is presently unaware of any documents sufficient to evidence any discussion with, disclosure to, or other manner of providing to a third party, or sale of or offer to sell, the inventions recited in the Asserted Claims of the Asserted Patents prior to the application dates or priority dates for the Asserted Patents. A diligent search continues for such documents and VirtaMove reserves the right to supplement this response.

G. Patent Rule 3-2(b) documents:

VirtaMove identifies the following non-privileged documents as related to evidencing conception and reduction to practice of each claimed invention of the Asserted Patents: VM_HPE_0000865–VM_HPE_0000880. A diligent search continues for additional documents and VirtaMove reserves the right to supplement this response.

H. Patent Rule 3-2(c) documents:

VirtaMove identifies the following documents as being the file histories for the Asserted Patents: VM_HPE_0000001–VM_HPE_0000864.

Dated: July 1, 2024

Respectfully submitted,

/s/ Reza Mirzaie

Reza Mirzaie
CA State Bar No. 246953
Marc A. Fenster
CA State Bar No. 181067
Neil A. Rubin
CA State Bar No. 250761
Amy E. Hayden
CA State Bar No. 287026
Jacob R. Buczko
CA State Bar No. 269408
James S. Tsuei
CA State Bar No. 285530
James A. Milkey
CA State Bar No. 281283
Christian W. Conkle
CA State Bar No. 306374
Jonathan Ma
CA State Bar No. 312773
Daniel Kolko (CA SBN 341680)
RUSS AUGUST & KABAT
12424 Wilshire Boulevard, 12th Floor
Los Angeles, CA 90025
Telephone: 310-826-7474
Email: rmirzaie@raklaw.com
Email: mfenster@raklaw.com
Email: nrubin@raklaw.com
Email: ahayden@raklaw.com
Email: jbuczko@raklaw.com
Email: jtsuei@raklaw.com
Email: jmilkey@raklaw.com
Email: cconkle@raklaw.com
Email: jma@raklaw.com
Email: dkolko@raklaw.com

Qi (Peter) Tong
4925 Greenville Ave., Suite 200
Dallas, TX 75206
Email: ptong@raklaw.com

**ATTORNEYS FOR PLAINTIFF
VIRTAMOVE, CORP.**

CERTIFICATE OF SERVICE

I certify that this document is being served upon counsel of record for Defendants
on July 1, 2024 via e-mail.

/s/ Reza Mirzaie
Reza Mirzaie

U.S. Patent No. 7,519,814 (“’814 Patent”)

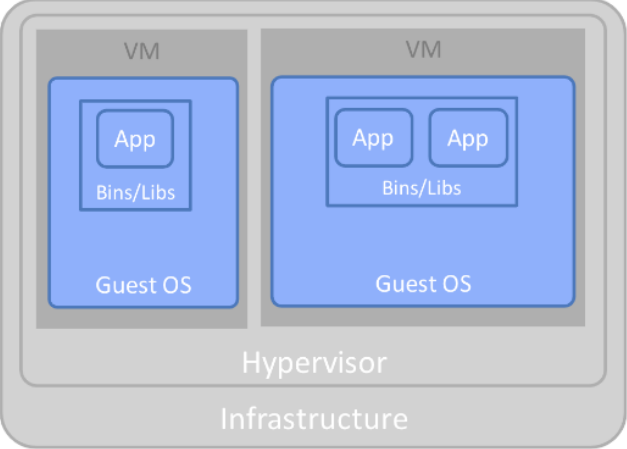
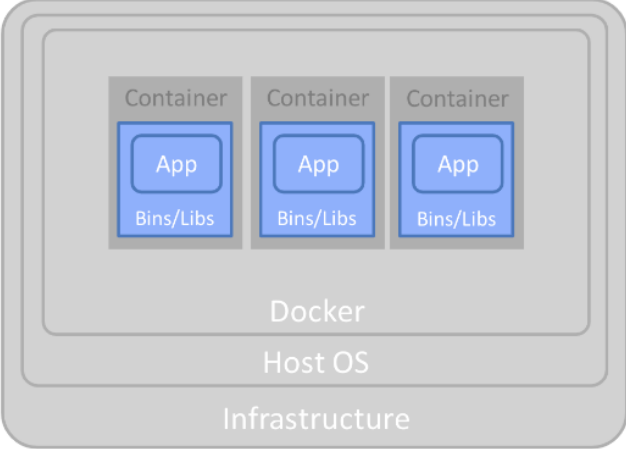
Accused Instrumentalities: IBM products and services using secure containerized applications, including without limitation IBM’s Cloud Kubernetes Service (IKS), IBM Cloud Private (ICP), and IBM Hybrid Cloud mesh, and all versions and variations thereof since the issuance of the asserted patent.

Each Accused Instrumentality infringes the claims in substantially the same way, and the evidence shown in this chart is similarly applicable to each Accused Instrumentality. Each claim limitation is literally infringed by each Accused Instrumentality. However, to the extent any claim limitation is not met literally, it is nonetheless met under the doctrine of equivalents because the differences between the claim limitation and each Accused Instrumentality would be insubstantial, and each Accused Instrumentality performs substantially the same function, in substantially the same way, to achieve the same result as the claimed invention. Notably, Defendant has not yet articulated which, if any, particular claim limitations it believes are not met by the Accused Instrumentalities.

Claim 1

Claim 1	Accused Instrumentalities
<p>[1pre] 1. In a system having a plurality of servers with operating systems that differ, operating in disparate computing environments, wherein each server includes a processor and an operating system including a kernel a set of associated local system files compatible with the processor, a method of providing at least some of the servers in the system with secure, executable, applications related to a service, wherein the applications are executed in a secure environment, wherein the applications each</p>	<p>To the extent the preamble is limiting, IBM practices, through the Accused Instrumentalities, in a system having a plurality of servers with operating systems that differ, operating in disparate computing environments, wherein each server includes a processor and an operating system including a kernel a set of associated local system files compatible with the processor, a method of providing at least some of the servers in the system with secure, executable, applications related to a service, wherein the applications are executed in a secure environment, wherein the applications each include an object executable by at least some of the different operating systems for performing a task related to the service, as claimed.</p> <p>For example, IBM Cloud Kubernetes Service runs on individual servers, each of which runs an independent operating system running either on bare metal, through an on-premises virtualized infrastructure, through one or more cloud services, or through any other supported deployment.</p> <p><i>See claim limitations below.</i></p> <p><i>See also, e.g.:</i></p>

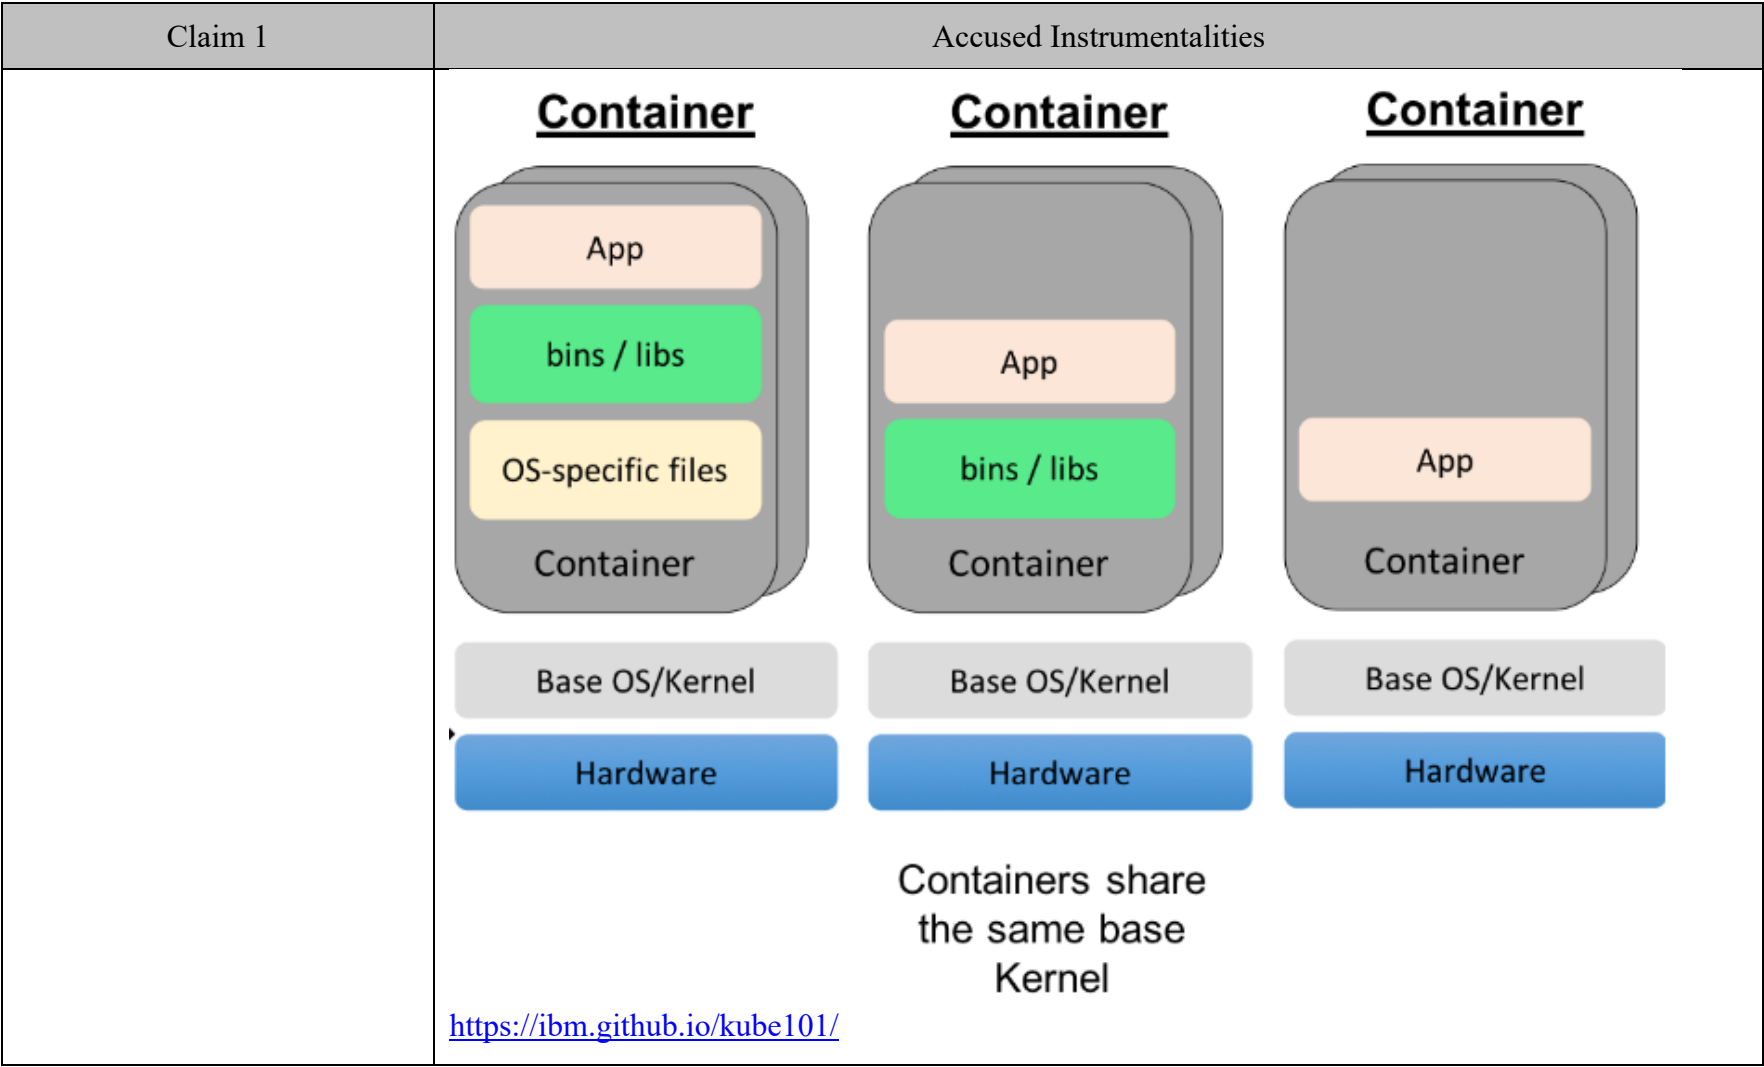
Claim 1	Accused Instrumentalities
<p>include an object executable by at least some of the different operating systems for performing a task related to the service, the method comprising:</p>	<p>IBM Cloud® Kubernetes Service provides a fully managed container service for Docker (OCI) containers, so clients can deploy containerized apps onto a pool of compute hosts and subsequently manage those containers. Containers are automatically scheduled and placed onto available compute hosts based on your requirements and availability in the cluster.</p> <p>https://www.ibm.com/products/kubernetes-service</p> <p>With IBM Cloud Kubernetes Service, you can deploy Docker containers into pods that run on your worker nodes. The worker nodes come with a set of add-on pods to help you manage your containers. Install more add-ons through Helm, a Kubernetes package manager. These add-ons can extend your apps with dashboards, logging, IBM Cloud and IBM Watson® services and more.</p> <p>https://www.ibm.com/products/kubernetes-service</p> <p>Docker is an open source platform that enables developers to build, deploy, run, update and manage <i>containers</i>—standardized, executable components that combine application source code with the operating system (OS) libraries and dependencies required to run that code in any environment.</p> <p>https://www.ibm.com/topics/docker</p>

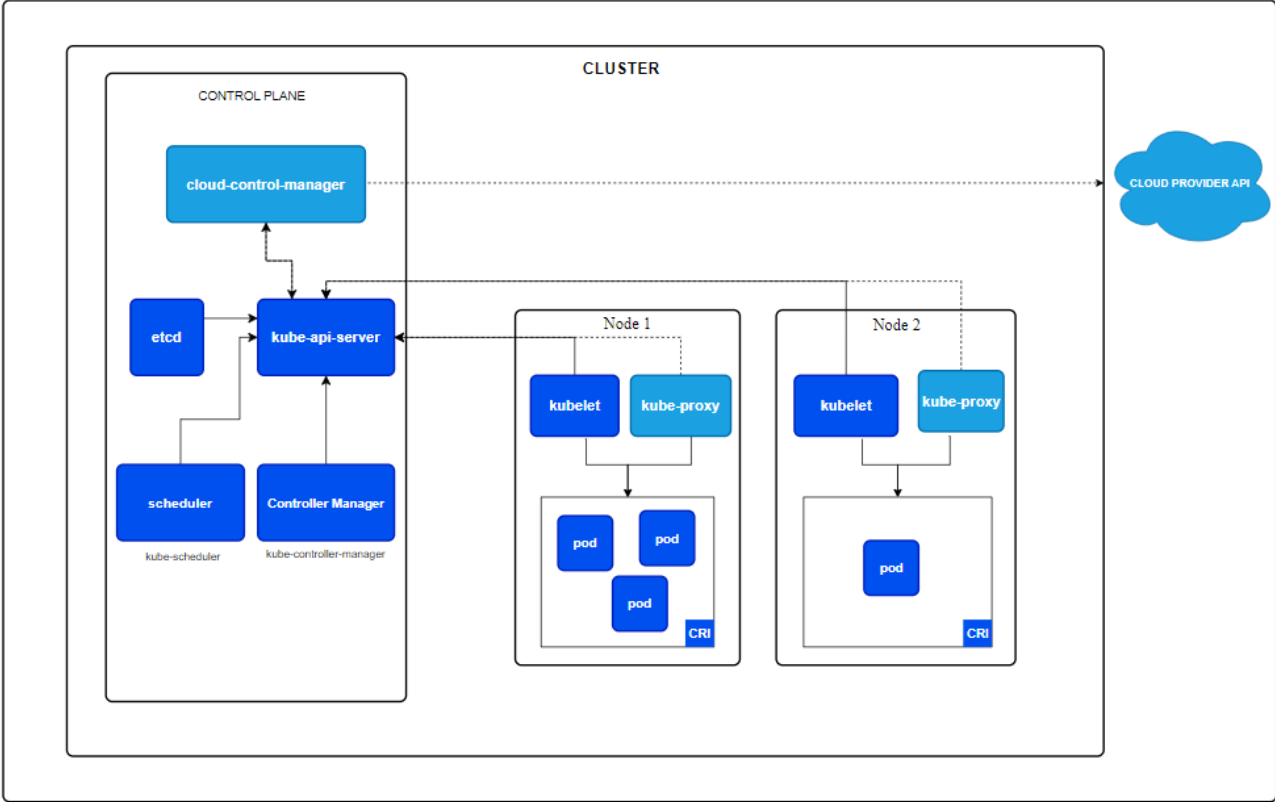
Claim 1	Accused Instrumentalities
	<div><div><p>Virtual Machines</p><p>The diagram illustrates the Virtual Machines architecture. At the base is the 'Infrastructure' layer. Above it is the 'Hypervisor' layer. The Hypervisor contains two 'VM' (Virtual Machine) boxes. Each VM box contains a 'Guest OS' layer. Inside each Guest OS, there is a 'Bins/Libs' layer. Within the Bins/Libs, there are 'App' (Application) boxes. The first VM contains one App box, and the second VM contains two App boxes.</p></div><div><p>Containers</p><p>The diagram illustrates the Containers architecture. At the base is the 'Infrastructure' layer. Above it is the 'Host OS' layer. Within the Host OS, there is a 'Docker' layer. The Docker layer contains three 'Container' boxes. Each Container box contains a 'Bins/Libs' layer. Within the Bins/Libs, there is an 'App' (Application) box. Each of the three Containers contains one App box.</p></div><p>https://developer.ibm.com/articles/true-benefits-of-moving-to-containers-1/</p></div>

Claim 1	Accused Instrumentalities
	<p>Containers are executable units of software in which application code is packaged along with its libraries and dependencies, in common ways so that the code can be run anywhere—whether it be on desktop, traditional IT or the cloud.</p> <p>To do this, containers take advantage of a form of operating system (OS) virtualization in which features of the OS kernel (e.g. Linux namespaces and cgroups, Windows silos and job objects) can be leveraged to isolate processes and control the amount of CPU, memory and disk that those processes can access.</p> <p>Containers are small, fast and portable because unlike a virtual machine, containers do not need to include a guest OS in every instance and can instead simply leverage the features and resources of the host OS.</p> <p>https://www.ibm.com/topics/containers</p> <p>With containers, you can isolate the ecosystem to run an application on any host OS (operating system). Containers can wrap code, runtimes, system tools, system libraries—everything that can be installed on a server. Containers are like virtual machines (VMs), but with a key difference in their architectural approach. Images that run on VMs have a full copy of the guest OS, including the necessary binaries and libraries. Images that run on containers share the OS kernel on the host.</p> <p>The Docker Engine builds and spins images on the containers. The engine is a lightweight container runtime that can run on almost any OS. You can run a container anywhere that a Docker Engine can be installed—on bare metal servers, clouds, and even inside a VM. You can move containers from one environment to another without recoding the application.</p> <p>Containers can help DevOps teams in three ways:</p> <ul style="list-style-type: none"> • Increase development productivity by reducing the time spent on environment setup • Eliminate issues that are caused by software dependencies • Avoid inconsistencies when applications are run in different environments <p>You can use IBM Cloud Kubernetes Service to run containers on IBM Cloud.</p> <p>https://www.ibm.com/garage/method/practices/run/tool_ibm_container/, last accessed on Nov. 17, 2023.</p>


Claim 1	Accused Instrumentalities
	<p>Containers use a form of operating system (OS) virtualization. Put simply, they leverage features of the host operating system to isolate processes and control the processes' access to CPUs, memory and disk space.</p> <p>https://www.ibm.com/blog/containers-vs-vms/</p> <p>Today Docker containerization also works with Microsoft Windows and Apple MacOS. Developers can run Docker containers on any operating system, and most leading cloud providers, including Amazon Web Services (AWS), Microsoft Azure, and IBM Cloud offer specific services to help developers build, deploy and run applications containerized with Docker.</p> <p>https://www.ibm.com/topics/docker</p>

Claim 1	Accused Instrumentalities
	<p>Containers are often referred to as “lightweight,” meaning they share the machine’s operating system kernel and do not require the overhead of associating an operating system within each application. Containers are inherently smaller in capacity than a VM and require less start-up time, allowing far more containers to run on the same compute capacity as a single VM. This drives higher server efficiencies and, in turn, reduces server and licensing costs.</p> <p>Containers encapsulate an application as a single executable package of software that bundles application code together with all of the related configuration files, libraries, and dependencies required for it to run. Containerized applications are “isolated” in that they do not bundle in a copy of the operating system. Instead, an open source runtime engine (such as the Docker runtime engine) is installed on the host’s operating system and becomes the conduit for containers to share an operating system with other containers on the same computing system.</p> <p>https://www.ibm.com/topics/containerization</p>








Claim 1	Accused Instrumentalities
	 <p>The diagram illustrates the Kubernetes cluster architecture. It is divided into a 'CONTROL PLANE' and a 'CLUSTER'. The 'CONTROL PLANE' contains components: 'cloud-control-manager', 'etcd', 'kube-api-server', 'scheduler' (labeled 'kube-scheduler'), and 'Controller Manager' (labeled 'kube-controller-manager'). The 'CLUSTER' contains 'Node 1' and 'Node 2'. Each node has a 'kubelet' and a 'kube-proxy'. 'Node 1' also contains three 'pod' boxes, and 'Node 2' contains one 'pod' box. A 'CNI' (Container Network Interface) label is present at the bottom of each node's pod area. A 'CLOUD PROVIDER API' cloud icon is connected to the 'cloud-control-manager' via a dashed line. Solid arrows show communication between 'etcd' and 'kube-api-server', and between 'kube-api-server' and the 'kubelet' on each node. Dashed arrows show communication between 'cloud-control-manager' and 'kube-api-server', and between 'kube-api-server' and the 'kube-proxy' on each node.</p> <p>Kubernetes cluster architecture https://kubernetes.io/docs/concepts/architecture/</p>
[1a] storing in memory accessible to at least some of the servers a plurality of secure containers of application software, each container	The method practiced by IBM through the Accused Instrumentalities includes a step of storing in memory accessible to at least some of the servers a plurality of secure containers of application software, each container comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications, for use with a local kernel residing permanently on one of the servers.

Claim 1	Accused Instrumentalities
<p>comprising one or more of the executable applications and a set of associated system files required to execute the one or more applications, for use with a local kernel residing permanently on one of the servers;</p>	<p>For example, IBM Cloud Kubernetes stores application containers, sometimes called Docker containers, container images, Kubernetes containers, or Kubernetes pods, in persistent storage available to each node running the application. The container might be in a format defined by the Open Container Initiative. This storage may be physically attached to the server or connected through any supported interconnect, including over a network. Each container includes the application software as well as a Linux user space required to execute the application, for example libc/glibc and other shared libraries, configuration files, etc. necessary for the application. For example, the container includes a base OS image, provided by IBM or by a third party, such as a CentOS, RHEL, or Ubuntu base image. The container is compatible with the host kernel, for example because the container libraries are linked against the Linux kernel, and the supported host operating systems also use the Linux kernel, which has a stable binary interface.</p> <p><i>See, e.g.:</i></p> <p>Containers use a form of operating system (OS) virtualization. Put simply, they leverage features of the host operating system to isolate processes and control the processes' access to CPUs, memory and disk space.</p> <p>https://www.ibm.com/blog/containers-vs-vms/</p> <p>Today Docker containerization also works with Microsoft Windows and Apple MacOS. Developers can run Docker containers on any operating system, and most leading cloud providers, including Amazon Web Services (AWS), Microsoft Azure, and IBM Cloud offer specific services to help developers build, deploy and run applications containerized with Docker.</p> <p>https://www.ibm.com/topics/docker</p>



Claim 1	Accused Instrumentalities
	<p data-bbox="667 272 982 321">Container images</p> <p data-bbox="667 349 1266 475">A <a data-bbox="667 349 835 373" href="https://kubernetes.io/docs/concepts/containers/">container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p data-bbox="634 500 1230 532"><a data-bbox="634 500 1230 532" href="https://kubernetes.io/docs/concepts/containers/">https://kubernetes.io/docs/concepts/containers/</p> <p data-bbox="659 581 1719 605">A Docker image is the basis for every container that you create with IBM Cloud® Kubernetes Service.</p> <p data-bbox="659 646 1890 751">An image is created from a Dockerfile, which is a file that contains instructions to build the image. A Dockerfile might reference build artifacts in its instructions that are stored separately, such as an app, the app's configuration, and its dependencies.</p> <p data-bbox="634 824 1449 857"><a data-bbox="634 824 1449 857" href="https://cloud.ibm.com/docs/containers?topic=containers-images">https://cloud.ibm.com/docs/containers?topic=containers-images</p>

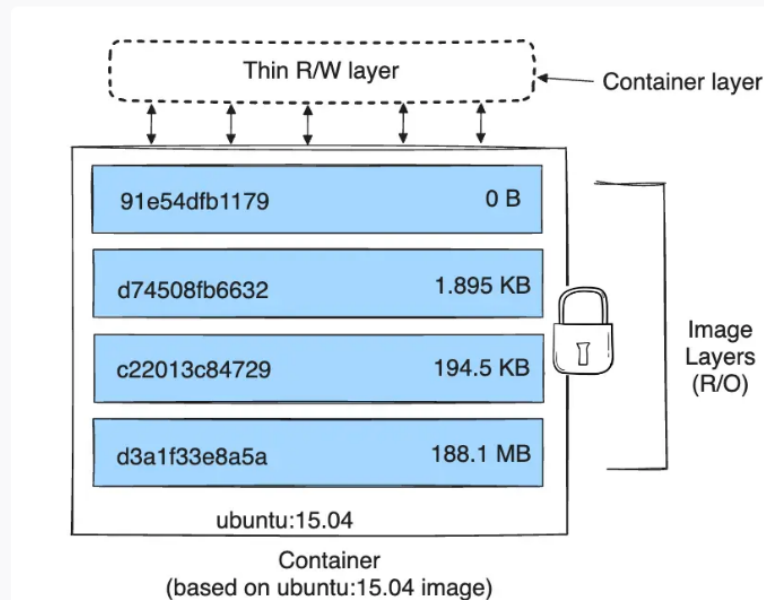
Claim 1	Accused Instrumentalities		
	Docker base image	Supported versions	Source of security notices
	Alpine	All stable versions with vendor security support.	Alpine SecDB database  .
	Debian	All stable versions with vendor security support. CVEs on binary packages that are associated with the Debian source package <code>linux</code> , such as <code>linux-libc-dev</code> , are not reported. Most of these binary packages are kernel and kernel modules, which are not run in container images.	Debian Security Bug Tracker  .
	GoogleContainerTools distroless	All stable versions with vendor security support.	GoogleContainerTools distroless  .
	Red Hat® Enterprise Linux® (RHEL)	RHEL/UBI 7, RHEL/UBI 8, and RHEL/UBI 9	Red Hat Security Data API  .
	Ubuntu	All stable versions with vendor security support.	Ubuntu CVE Tracker  .
	https://cloud.ibm.com/docs/Registry?topic=Registry-va_index&interface=ui		

Claim 1	Accused Instrumentalities
	<h2 data-bbox="646 277 1272 342">About storage drivers</h2> <p data-bbox="646 391 1871 516">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="646 581 1564 646">Storage drivers versus Docker volumes</h2> <p data-bbox="646 678 1913 943">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="646 992 1906 1117">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1339 1040 1535 1068" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="636 1149 1226 1182"><a data-bbox="636 1149 1226 1182" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="657 272 1083 329">Images and layers</h2> <p data-bbox="657 367 1822 443">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="674 513 1453 824"> # syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="657 889 1898 1198">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="636 1219 1226 1252">https://docs.docker.com/storage/storagedriver/</p>

Each layer is only a set of differences from the layer before it. Note that both *adding*, and *removing* files will result in a new layer. In the example above, the `$HOME/.cache` directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the [Best practices for writing Dockerfiles](#) and [use multi-stage builds](#) sections to learn how to optimize your Dockerfiles for efficient images.

The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an `ubuntu:15.04` image.



<https://docs.docker.com/storage/storagedriver/>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 289 919 347">Volumes</h2> <p data-bbox="653 402 1906 532">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="634 557 1308 589">https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2 data-bbox="653 639 1224 688">Container environment</h2> <p data-bbox="653 727 1474 792">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="695 829 1451 992" style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster. <p data-bbox="634 1024 1528 1057">https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="659 280 877 342">Images</h2> <p data-bbox="659 375 1522 526">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="659 565 1528 634">You typically create a container image of your application and push it to a registry before referring to it in a <code>Pod</code>.</p> <p data-bbox="634 662 1329 695">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="653 740 919 795">Volumes</h2> <p data-bbox="653 837 1528 1279">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a <code>Pod</code> and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes <code>volume</code> abstraction solves both of these problems. Familiarity with <code>Pods</code> is suggested.</p> <p data-bbox="634 1307 1308 1339">https://kubernetes.io/docs/concepts/storage/volumes/</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="674 282 1295 342"><h2>Open Container Initiative</h2><hr/></div> <div data-bbox="674 404 1182 451"><h3>Image Format Specification</h3><hr/></div> <div data-bbox="674 498 1890 573"><p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p></div> <div data-bbox="674 609 1900 683"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="634 709 1478 781"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>


Claim 1	Accused Instrumentalities
	<div data-bbox="648 277 831 318">Overview</div> <div data-bbox="648 371 1892 613"><p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p></div> <div data-bbox="648 656 1908 1036"><div data-bbox="669 792 982 881"><pre>public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } }</pre></div><div data-bbox="1003 824 1052 846">→</div><div data-bbox="1075 662 1312 987"><div data-bbox="1075 662 1123 711"></div><div data-bbox="1123 662 1312 987">/bin/java /opt/app.jar /lib/libc</div></div><div data-bbox="1171 1008 1234 1036">layer</div><div data-bbox="1333 833 1360 854">+</div><div data-bbox="1360 662 1617 987"><div data-bbox="1360 662 1409 711"></div><div data-bbox="1409 662 1617 987">{ "manifests": { "platform": { "os": "linux", ... } } }</div></div><div data-bbox="1430 1008 1577 1036">image index</div><div data-bbox="1627 833 1654 854">+</div><div data-bbox="1654 662 1908 987"><div data-bbox="1654 662 1703 711"></div><div data-bbox="1703 662 1908 987">{ ... "config": { "Cmd": ["java", "-jar", "app.jar"], ... } }</div></div><div data-bbox="1759 1008 1833 1036">config</div></div> <div data-bbox="632 1065 1478 1133"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 272 1297 334">OCI Image Configuration</h2> <p data-bbox="653 386 1913 548">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p data-bbox="653 586 1661 618">This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p data-bbox="632 651 1503 724">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>






Claim 1	Accused Instrumentalities
	<p>Layer</p> <ul style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p>Image JSON</p> <ul style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<ul style="list-style-type: none"> • rootfs object, REQUIRED <p>The rootfs key references the layer content addresses used by the image. This makes the image config hash depend on the filesystem hash.</p> <ul style="list-style-type: none"> ◦ type string, REQUIRED <p>MUST be set to <code>layers</code>. Implementations MUST generate an error if they encounter a unknown value while verifying or unpacking an image.</p> <ul style="list-style-type: none"> ◦ diff_ids array of strings, REQUIRED <p>An array of layer content hashes (<code>DiffIDs</code>), in order from first to last.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>
<p>[1b] wherein the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems,</p>	<p>In the method practiced by IBM through the Accused Instrumentalities, the set of associated system files are compatible with a local kernel of at least some of the plurality of different operating systems.</p> <p>The system files in the container are compatible with the host kernel, for example because they are linked against the Linux kernel and the supported host operating systems also use the Linux kernel, which has a stable binary interface.</p> <p><i>See, e.g.:</i></p>

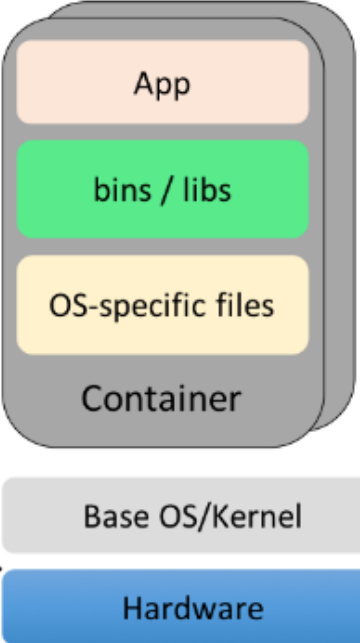
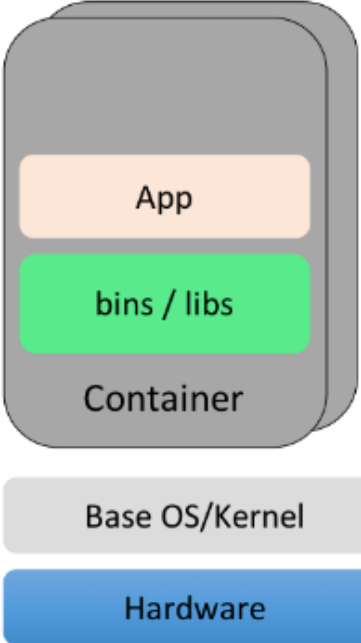
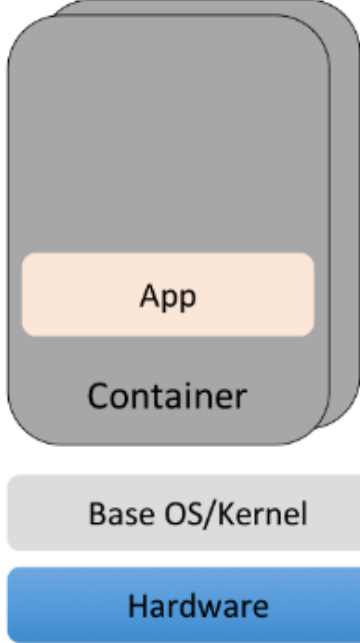
Claim 1	Accused Instrumentalities
	<p>Containers are often referred to as “lightweight,” meaning they share the machine’s operating system kernel and do not require the overhead of associating an operating system within each application. Containers are inherently smaller in capacity than a VM and require less start-up time, allowing far more containers to run on the same compute capacity as a single VM. This drives higher server efficiencies and, in turn, reduces server and licensing costs.</p> <p>Containers encapsulate an application as a single executable package of software that bundles application code together with all of the related configuration files, libraries, and dependencies required for it to run. Containerized applications are “isolated” in that they do not bundle in a copy of the operating system. Instead, an open source runtime engine (such as the Docker runtime engine) is installed on the host’s operating system and becomes the conduit for containers to share an operating system with other containers on the same computing system.</p> <p>https://www.ibm.com/topics/containerization</p>



Claim 1	Accused Instrumentalities
	<p data-bbox="661 284 1717 316">A Docker image is the basis for every container that you create with IBM Cloud® Kubernetes Service.</p> <p data-bbox="661 349 1890 462">An image is created from a Dockerfile, which is a file that contains instructions to build the image. A Dockerfile might reference build artifacts in its instructions that are stored separately, such as an app, the app's configuration, and its dependencies.</p> <p data-bbox="634 527 1449 560">https://cloud.ibm.com/docs/containers?topic=containers-images</p>

Claim 1	Accused Instrumentalities		
	Docker base image	Supported versions	Source of security notices
	Alpine	All stable versions with vendor security support.	Alpine SecDB database  .
	Debian	All stable versions with vendor security support. CVEs on binary packages that are associated with the Debian source package <code>linux</code> , such as <code>linux-libc-dev</code> , are not reported. Most of these binary packages are kernel and kernel modules, which are not run in container images.	Debian Security Bug Tracker  .
	GoogleContainerTools distroless	All stable versions with vendor security support.	GoogleContainerTools distroless  .
	Red Hat® Enterprise Linux® (RHEL)	RHEL/UBI 7, RHEL/UBI 8, and RHEL/UBI 9	Red Hat Security Data API  .
	Ubuntu	All stable versions with vendor security support.	Ubuntu CVE Tracker  .
	https://cloud.ibm.com/docs/Registry?topic=Registry-va_index&interface=ui		
[1c] the containers of application software excluding a kernel,	In the method practiced by IBM through the Accused Instrumentalities, the containers of application software exclude a kernel. <i>See, e.g.:</i>		

Claim 1	Accused Instrumentalities
	<p>Containers are often referred to as “lightweight,” meaning they share the machine’s operating system kernel and do not require the overhead of associating an operating system within each application. Containers are inherently smaller in capacity than a VM and require less start-up time, allowing far more containers to run on the same compute capacity as a single VM. This drives higher server efficiencies and, in turn, reduces server and licensing costs.</p> <p>https://www.ibm.com/topics/containerization</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="640 272 1843 1182"><div><p><u>Container</u></p><p>App</p><p>bins / libs</p><p>OS-specific files</p><p>Container</p><p>Base OS/Kernel</p><p>Hardware</p></div><div><p><u>Container</u></p><p>App</p><p>bins / libs</p><p>Container</p><p>Base OS/Kernel</p><p>Hardware</p></div><div><p><u>Container</u></p><p>App</p><p>Container</p><p>Base OS/Kernel</p><p>Hardware</p></div><p>Containers share the same base Kernel</p><p>https://ibm.github.io/kube101/</p></div>
[1d] wherein some or all of the associated system files within a container stored in memory are utilized in place of the	In the method practiced by IBM through the Accused Instrumentalities, some or all of the associated system files within a container stored in memory are utilized in place of the associated local system files that remain resident on the server.

Claim 1	Accused Instrumentalities
<p>associated local system files that remain resident on the server,</p>	<p>For example, each container will utilize its own local system files, including libraries such as libc/glibc and configuration files, not the corresponding libraries and configuration files of the host OS.</p> <p><i>See, e.g.:</i></p> <p>Rather than spinning up an entire virtual machine, containerization packages together everything needed to run a single application or microservice (along with runtime libraries they need to run). The container includes all the code, its dependencies and even the operating system itself. This enables applications to run almost anywhere — a desktop computer, a traditional IT infrastructure or the cloud.</p> <p>Containers use a form of operating system (OS) virtualization. Put simply, they leverage features of the host operating system to isolate processes and control the processes' access to CPUs, memory and desk space.</p> <p>https://www.ibm.com/blog/containers-vs-vms/</p>
<p>[1e] wherein said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server,</p>	<p>In the method practiced by IBM through the Accused Instrumentalities, said associated system files utilized in place of the associated local system files are copies or modified copies of the associated local system files that remain resident on the server.</p> <p>For example, in some cases the host OS and container will use one or more identical system files, for example when both the host and the container incorporate the same Linux distribution version, or when both host and container use the same version of libc. In other cases modified copies are used instead, for example when different versions of the same library, or configuration files with different parameters, are used by the host and container.</p> <p><i>See, e.g.:</i></p>

Claim 1	Accused Instrumentalities
	<p>Containerization is the packaging of software code with just the operating system (OS) libraries and dependencies required to run the code to create a single lightweight executable—called a container—that runs consistently on any infrastructure. More portable and resource-efficient than virtual machines (VMs), containers have become the de facto compute units of modern cloud-native applications.</p> <p>Containerization allows developers to create and deploy applications faster and more securely. With traditional methods, code is developed in a specific computing environment which, when transferred to a new location, often results in bugs and errors. For example, when a developer transfers code from a desktop computer to a VM or from a Linux to a Windows operating system. Containerization eliminates this problem by bundling the application code together with the related configuration files, libraries, and dependencies required for it to run. This single package of software or “container” is abstracted away from the host operating system, and hence, it stands alone and becomes portable—able to run across any platform or cloud, free of issues.</p> <p>https://www.ibm.com/topics/containerization</p>

Claim 1	Accused Instrumentalities
	<p>With containers, you can isolate the ecosystem to run an application on any host OS (operating system). Containers can wrap code, runtimes, system tools, system libraries—everything that can be installed on a server. Containers are like virtual machines (VMs), but with a key difference in their architectural approach. Images that run on VMs have a full copy of the guest OS, including the necessary binaries and libraries. Images that run on containers share the OS kernel on the host.</p> <p>The Docker Engine builds and spins images on the containers. The engine is a lightweight container runtime that can run on almost any OS. You can run a container anywhere that a Docker Engine can be installed—on bare metal servers, clouds, and even inside a VM. You can move containers from one environment to another without recoding the application.</p> <p>Containers can help DevOps teams in three ways:</p> <ul style="list-style-type: none"> • Increase development productivity by reducing the time spent on environment setup • Eliminate issues that are caused by software dependencies • Avoid inconsistencies when applications are run in different environments <p>You can use IBM Cloud Kubernetes Service to run containers on IBM Cloud.</p> <p>https://www.ibm.com/garage/method/practices/run/tool_ibm_container/, last accessed on Nov. 17, 2023.</p>
<p>[1f] and wherein the application software cannot be shared between the plurality of secure containers of application software,</p>	<p>In the method practiced by IBM through the Accused Instrumentalities, the application software cannot be shared between the plurality of secure containers of application software.</p> <p>For example, each container has an isolated runtime environment that cannot be accessed by other containers, for example including a per-container writeable layer or other ephemeral per-container storage. For another example, when the plurality of secure containers each corresponds to a different container image, each container cannot access another container's image and therefore application software.</p> <p><i>See, e.g.:</i></p> <p>Containers are made possible by process isolation and virtualization capabilities built into the Linux kernel. These capabilities—such as <i>control groups</i> (Cgroups) for allocating resources among processes, and <i>namespaces</i> for restricting a processes access or visibility into other resources or areas of the system—enable multiple application components to share the resources of a single instance of the host operating system in much the same way that a hypervisor enables multiple virtual machines (VMs) to share the CPU, memory and other resources of a single hardware server.</p> <p>https://www.ibm.com/topics/docker</p>

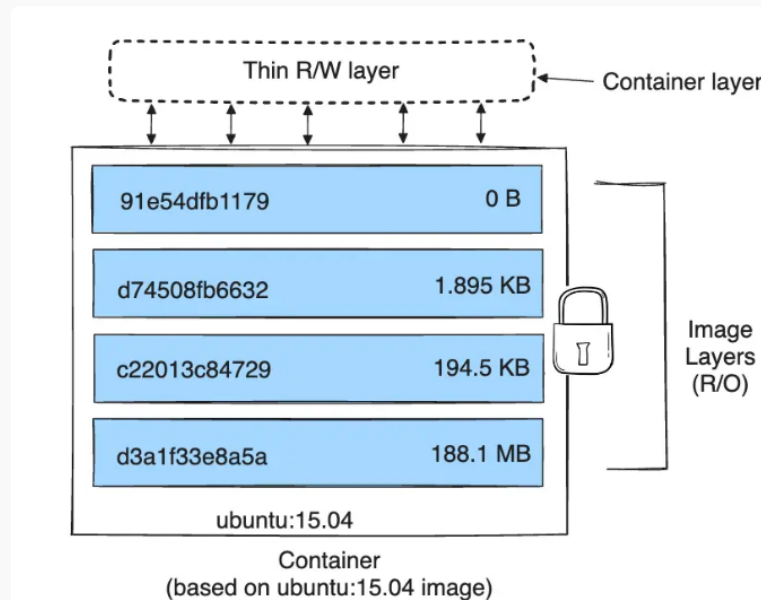
Claim 1	Accused Instrumentalities
	<p>Fault isolation: Each containerized application is isolated and operates independently of others. The failure of one container does not affect the continued operation of any other containers. Development teams can identify and correct any technical issues within one container without any downtime in other containers. Also, the container engine can leverage any OS security isolation techniques—such as SELinux access control—to isolate faults within containers.</p> <p>https://www.ibm.com/topics/containerization</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="642 277 1272 342">About storage drivers</h2> <p data-bbox="642 391 1871 513">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="642 583 1564 639">Storage drivers versus Docker volumes</h2> <p data-bbox="642 678 1913 943">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="642 995 1902 1117">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1339 1044 1530 1068" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="642 1149 1224 1182"><a data-bbox="642 1149 1224 1182" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="657 272 1081 329">Images and layers</h2> <p data-bbox="657 367 1822 443">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="674 513 1451 824"> # syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="657 889 1900 1198">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="634 1219 1224 1252">https://docs.docker.com/storage/storagedriver/</p>

Each layer is only a set of differences from the layer before it. Note that both *adding*, and *removing* files will result in a new layer. In the example above, the `$HOME/.cache` directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the [Best practices for writing Dockerfiles](#) and [use multi-stage builds](#) sections to learn how to optimize your Dockerfiles for efficient images.

The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an `ubuntu:15.04` image.



<https://docs.docker.com/storage/storagedriver/>

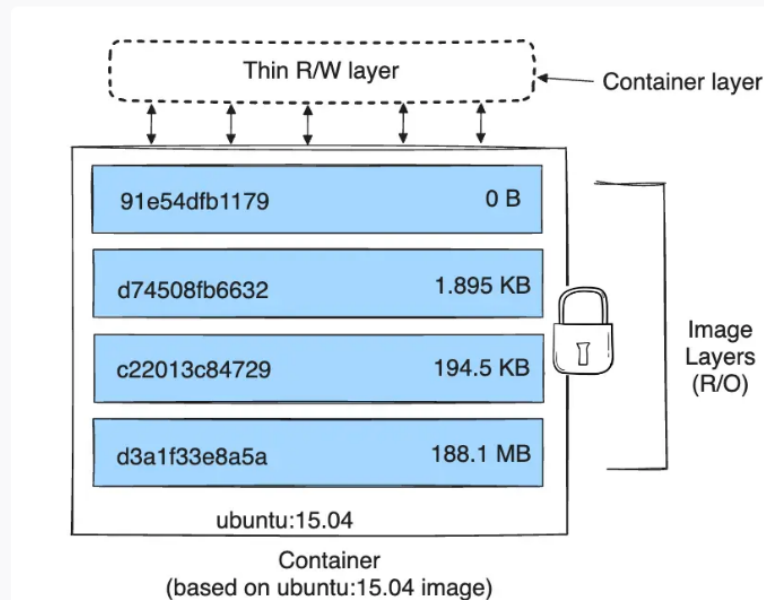
Claim 1	Accused Instrumentalities
<p>[1g] and wherein each of the containers has a unique root file system that is different from an operating system's root file system.</p>	<p>In the method practiced by IBM through the Accused Instrumentalities, each of the containers has a unique root file system that is different from an operating system's root file system.</p> <p>For example, the container's root file system comprises the image layer(s), an ephemeral writeable layer (e.g., in Docker terminology the container layer), and optionally one or more volumes. This root file system is distinct and isolated from the host operating system's root file system.</p> <p><i>See, e.g.:</i></p> <p>Limit the number of privileged containers. Containers run as a separate Linux process on the compute host that is isolated from other processes. Although users have root access inside the container, the permissions of this user are limited outside the container to protect other Linux processes, the host file system, and host devices. Some apps require access to the host file system or advanced permissions to run properly. You can run containers in privileged mode to allow the container the same access as the processes running on the compute host. Keep in mind that privileged containers can cause huge damage to the cluster and the underlying compute host if they become compromised. Try to limit the number of containers that run in privileged mode and consider changing the configuration for your app so that the app can run without advanced permissions.</p> <p>https://cloud.ibm.com/docs/containers?topic=containers-security</p> <p>Containers are made possible by process isolation and virtualization capabilities built into the Linux kernel. These capabilities—such as <i>control groups</i> (Cgroups) for allocating resources among processes, and <i>namespaces</i> for restricting a processes access or visibility into other resources or areas of the system—enable multiple application components to share the resources of a single instance of the host operating system in much the same way that a hypervisor enables multiple <i>virtual machines</i> (VMs) to share the CPU, memory and other resources of a single hardware server.</p> <p>https://www.ibm.com/topics/docker</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="646 277 1272 342">About storage drivers</h2> <p data-bbox="646 391 1871 516">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="646 581 1564 646">Storage drivers versus Docker volumes</h2> <p data-bbox="646 678 1913 943">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="646 992 1906 1117">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1339 1040 1535 1068" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="636 1149 1224 1182"><a data-bbox="636 1149 1224 1182" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="657 272 1081 329">Images and layers</h2> <p data-bbox="657 367 1822 443">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="674 513 1451 824"> # syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="657 889 1900 1198">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="636 1219 1224 1252">https://docs.docker.com/storage/storagedriver/</p>

Each layer is only a set of differences from the layer before it. Note that both *adding*, and *removing* files will result in a new layer. In the example above, the `$HOME/.cache` directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the [Best practices for writing Dockerfiles](#) and [use multi-stage builds](#) sections to learn how to optimize your Dockerfiles for efficient images.

The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an `ubuntu:15.04` image.



<https://docs.docker.com/storage/storagedriver/>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 289 919 349">Volumes</h2> <p data-bbox="653 402 1906 532">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="634 557 1308 589">https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2 data-bbox="653 638 1226 690">Container environment</h2> <p data-bbox="653 727 1474 792">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="695 829 1451 992" style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster. <p data-bbox="634 1024 1528 1057">https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="659 280 877 342">Images</h2> <p data-bbox="659 375 1520 526">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="659 565 1528 634">You typically create a container image of your application and push it to a registry before referring to it in a <code>Pod</code>.</p> <p data-bbox="634 662 1329 695">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="653 740 919 802">Volumes</h2> <p data-bbox="653 837 1528 1279">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers. One problem occurs when a container crashes or is stopped. Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a <code>Pod</code> and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes <code>volume</code> abstraction solves both of these problems. Familiarity with <code>Pods</code> is suggested.</p> <p data-bbox="634 1307 1308 1339">https://kubernetes.io/docs/concepts/storage/volumes/</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="674 282 1295 342"><h2>Open Container Initiative</h2><hr/></div> <div data-bbox="674 406 1182 451"><h3>Image Format Specification</h3><hr/></div> <div data-bbox="674 498 1890 573"><p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p></div> <div data-bbox="674 609 1900 683"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="634 709 1478 781"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 1	Accused Instrumentalities
	<div data-bbox="648 277 831 318">Overview</div> <div data-bbox="648 371 1892 615"><p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p></div> <div data-bbox="648 656 1908 1036"><div data-bbox="669 792 982 881"><pre>public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } }</pre></div><div data-bbox="1003 824 1052 846">→</div><div data-bbox="1073 656 1312 1036"><div data-bbox="1073 656 1121 711"></div><div data-bbox="1094 656 1312 987"><p>/bin/java /opt/app.jar /lib/libc</p></div><div data-bbox="1171 1003 1234 1036">layer</div></div><div data-bbox="1333 833 1360 854">+</div><div data-bbox="1360 656 1617 1036"><div data-bbox="1360 656 1409 711"></div><div data-bbox="1381 656 1617 987"><pre>{ "manifests": { "platform": { "os": "linux", ... } } }</pre></div><div data-bbox="1430 1003 1577 1036">image index</div></div><div data-bbox="1627 833 1654 854">+</div><div data-bbox="1654 656 1908 1036"><div data-bbox="1654 656 1703 711"></div><div data-bbox="1675 656 1908 987"><pre>{ ... "config": { "Cmd": ["java", "-jar", "app.jar"], ... } }</pre></div><div data-bbox="1759 1003 1833 1036">config</div></div></div> <div data-bbox="632 1063 1480 1133"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 1	Accused Instrumentalities
	<div data-bbox="653 272 1297 332"><h2>OCI Image Configuration</h2></div> <div data-bbox="653 386 1915 548"><p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p><p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p></div>

Claim 1	Accused Instrumentalities
	<p>Layer</p> <ul style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p>Image JSON</p> <ul style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<ul style="list-style-type: none"> • rootfs object, REQUIRED The rootfs key references the layer content addresses used by the image. This makes the image config hash depend on the filesystem hash. ◦ type string, REQUIRED MUST be set to <code>layers</code>. Implementations MUST generate an error if they encounter a unknown value while verifying or unpacking an image. ◦ diff_ids array of strings, REQUIRED An array of layer content hashes (<code>DiffIDs</code>), in order from first to last. https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md

Claim 2

Claim 2	Accused Instrumentalities
2. A method as defined in claim 1, wherein each container has an execution file associated therewith for starting the one or more applications.	<p>IBM practices, through the Accused Instrumentalities, a method as defined in claim 1, wherein each container has an execution file associated therewith for starting the one or more applications.</p> <p>For example, a container image has an associated image configuration comprising information for starting the one or more applications. This can be an Open Containers Initiative image configuration.</p> <p><i>See, e.g.:</i></p>

Claim 2	Accused Instrumentalities
	<div data-bbox="625 256 1251 315"><h2>Open Container Initiative</h2><hr/></div> <div data-bbox="625 378 1136 423"><h3>Image Format Specification</h3><hr/></div> <div data-bbox="625 470 1848 545"><p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p></div> <div data-bbox="625 579 1854 654"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="585 682 1432 751"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 2	Accused Instrumentalities
	<div data-bbox="604 248 785 289"><h3>Overview</h3></div> <div data-bbox="604 345 1848 586"><p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p></div> <div data-bbox="604 630 1864 1008"><p>The diagram illustrates the components of an OCI image. On the left, a code block shows a Java class: <pre>public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } }</pre>. An arrow points from this code to a cylinder labeled 'layer' containing the paths <code>/bin/java</code>, <code>/opt/app.jar</code>, and <code>/lib/libc</code>. To the right of the layer is a plus sign, followed by a document icon labeled 'image index' containing a JSON snippet: <pre>{ "manifests": { "platform": { "os": "linux", ... } } }</pre>. Another plus sign follows, leading to a document icon labeled 'config' containing a JSON snippet: <pre>{ ... "config": { "Cmd": ["java", "-jar", "app.jar"], ... } }</pre></p></div> <div data-bbox="588 1036 1432 1105"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

Claim 2	Accused Instrumentalities
	<div data-bbox="604 245 1251 305"><h2>OCI Image Configuration</h2></div> <div data-bbox="604 358 1871 521"><p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p></div> <div data-bbox="604 558 1612 591"><p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p></div> <div data-bbox="588 623 1457 695"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p></div>

Claim 2	Accused Instrumentalities
	<ul style="list-style-type: none"> • config object, OPTIONAL <p>The execution parameters which SHOULD be used as a base when running a container using the image. This field can be <code>null</code>, in which case any execution parameters should be specified at creation of the container.</p> <ul style="list-style-type: none"> ◦ Env array of strings, OPTIONAL <p>Entries are in the format of <code>VARNAME=VARVALUE</code>. These values act as defaults and are merged with any specified when creating a container.</p> <ul style="list-style-type: none"> ◦ Entrypoint array of strings, OPTIONAL <p>A list of arguments to use as the command to execute when the container starts. These values act as defaults and may be replaced by an entrypoint specified when creating a container.</p> <ul style="list-style-type: none"> ◦ Cmd array of strings, OPTIONAL <p>Default arguments to the entrypoint of the container. These values act as defaults and may be replaced by any specified when creating a container. If an <code>Entrypoint</code> value is not specified, then the first entry of the <code>Cmd</code> array SHOULD be interpreted as the executable to run.</p> <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 6

Claim 6	Accused Instrumentalities
<p>6. A method as defined in claim 2, comprising the step of assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address.</p>	<p>IBM practices, through the Accused Instrumentalities, a method as defined in claim 2, comprising the step of assigning a unique associated identity to each of a plurality of the containers, wherein the identity includes at least one of IP address, host name, and MAC address.</p> <p>For example, Kubernetes containers have an associated hostname, which in the case of a single-container Pod is the unique identity of that container. For another example, Kubernetes pods have an associated hostname, which is unique. For another example, a networked Kubernetes pod has an assigned IPv4 and/or IPv6 address. For another example, a Docker container has an IP address and a hostname.</p> <p><i>See, e.g.:</i></p> <p>Container information</p> <p>The <i>hostname</i> of a Container is the name of the Pod in which the Container is running. It is available through the <code>hostname</code> command or the <code>gethostname</code> function call in libc.</p> <p>The Pod name and namespace are available as environment variables through the downward API.</p> <p>User defined environment variables from the Pod definition are also available to the Container, as are any environment variables specified statically in the container image.</p> <p>https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 6	Accused Instrumentalities
	<p data-bbox="604 245 1201 293">IP address and hostname</p> <p data-bbox="604 337 1860 461">By default, the container gets an IP address for every Docker network it attaches to. A container receives an IP address out of the IP subnet of the network. The Docker daemon performs dynamic subnetting and IP address allocation for containers. Each network also has a default subnet mask and gateway.</p> <p data-bbox="604 511 1843 678">You can connect a running container to multiple networks, either by passing the <code>--network</code> flag multiple times when creating the container, or using the <code>docker network connect</code> command for already running containers. In both cases, you can use the <code>--ip</code> or <code>--ip6</code> flags to specify the container's IP address on that particular network.</p> <p data-bbox="604 732 1850 855">In the same way, a container's hostname defaults to be the container's ID in Docker. You can override the hostname using <code>--hostname</code>. When connecting to an existing network using <code>docker network connect</code>, you can use the <code>--alias</code> flag to specify an additional network alias for the container on that network.</p> <p data-bbox="588 889 1014 917">https://docs.docker.com/network/</p>

Claim 9

Claim 9	Accused Instrumentalities
<p data-bbox="203 1078 682 1360">9. A method as defined in claim 2, wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed.</p>	<p data-bbox="714 1078 1850 1219">IBM practices, through the Accused Instrumentalities, a method as defined in claim 2, wherein server information related to hardware resource usage including at least one of CPU memory, network bandwidth, and disk allocation is associated with at least some of the containers prior to the applications within the containers being executed.</p> <p data-bbox="714 1243 1843 1351">For example, Kubernetes tracks and limits resource usage, including CPU and memory resources. For another example, Docker tracks and limits resource usage, including CPU and memory resources.</p> <p data-bbox="714 1375 831 1409"><i>See, e.g.:</i></p>

Resource Management for Pods and Containers

When you specify a Pod, you can optionally specify how much of each resource a container needs. The most common resources to specify are CPU and memory (RAM); there are others.

When you specify the resource *request* for containers in a Pod, the kube-scheduler uses this information to decide which node to place the Pod on. When you specify a resource *limit* for a container, the kubelet enforces those limits so that the running container is not allowed to use more of that resource than the limit you set. The kubelet also reserves at least the *request* amount of that system resource specifically for that container to use.

Requests and limits

If the node where a Pod is running has enough of a resource available, it's possible (and allowed) for a container to use more resource than its *request* for that resource specifies. However, a container is not allowed to use more than its resource *limit*.

For example, if you set a *memory request* of 256 MiB for a container, and that container is in a Pod scheduled to a Node with 8GiB of memory and no other Pods, then the container can try to use more RAM.

If you set a *memory limit* of 4GiB for that container, the kubelet (and container runtime) enforce the limit. The runtime prevents the container from using more than the configured resource limit. For example: when a process in the container tries to consume more than

Claim 9	Accused Instrumentalities
	<p>the allowed amount of memory, the system kernel terminates the process that attempted the allocation, with an out of memory (OOM) error.</p> <p>Limits can be implemented either reactively (the system intervenes once it sees a violation) or by enforcement (the system prevents the container from ever exceeding the limit). Different runtimes can have different ways to implement the same restrictions.</p> <p>https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</p> <p>Runtime options with Memory, CPUs, and GPUs</p> <p>By default, a container has no resource constraints and can use as much of a given resource as the host's kernel scheduler allows. Docker provides ways to control how much memory, or CPU a container can use, setting runtime configuration flags of the <code>docker run</code> command. This section provides details on when you should set such limits and the possible implications of setting them.</p> <p>Limit a container's access to memory</p> <p>Docker can enforce hard or soft memory limits.</p> <ul style="list-style-type: none"> • Hard limits lets the container use no more than a fixed amount of memory. • Soft limits lets the container use as much memory as it needs unless certain conditions are met, such as when the kernel detects low memory or contention on the host machine. <p>https://docs.docker.com/config/containers/resource_constraints/</p>

Claim 10

Claim 10	Accused Instrumentalities
<p>10. A method as defined in claim 2, wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof.</p>	<p>IBM practices, through the Accused Instrumentalities, a method as defined in claim 2, wherein in operation when an application residing within a container is executed, said application has no access to system files or applications in other containers or to system files within the operating system during execution thereof.</p> <p><i>See, e.g.:</i></p> <p>Containers are made possible by process isolation and virtualization capabilities built into the Linux kernel. These capabilities—such as <i>control groups</i> (Cgroups) for allocating resources among processes, and <i>namespaces</i> for restricting a processes access or visibility into other resources or areas of the system—enable multiple application components to share the resources of a single instance of the host operating system in much the same way that a hypervisor enables multiple virtual machines (VMs) to share the CPU, memory and other resources of a single hardware server.</p> <p>https://www.ibm.com/topics/docker</p> <p>Fault isolation: Each containerized application is isolated and operates independently of others. The failure of one container does not affect the continued operation of any other containers. Development teams can identify and correct any technical issues within one container without any downtime in other containers. Also, the container engine can leverage any OS security isolation techniques—such as SELinux access control—to isolate faults within containers.</p> <p>https://www.ibm.com/topics/containerization</p>

Claim 31

<u>Claim 31</u>	<u>Accused Instrumentalities</u>
<p><u>[31pre] A computing system for performing a plurality of tasks each comprising a plurality of processes comprising:</u></p>	<p><u>To the extent the preamble is construed as a limitation, each Accused Instrumentality is or comprises a computing system for performing a plurality of tasks each comprising a plurality of processes.</u></p> <p><u>See claim limitations below. See also analysis and evidence for [1pre] above.</u></p>
<p><u>[31a] a system having a plurality of secure containers of associated files accessible to, and for execution on, one or more servers, each container being mutually exclusive of the other, such that read/write files within a container cannot be shared with other containers, each container of files is said to have its own unique identity associated therewith, said identity comprising at least one of an IP address, a host name, and a Mac address</u></p>	<p><u>Each Accused Instrumentality comprises a system having a plurality of secure containers of associated files accessible to, and for execution on, one or more servers, each container being mutually exclusive of the other, such that read/write files within a container cannot be shared with other containers, each container of files is said to have its own unique identity associated therewith, said identity comprising at least one of an IP address, a host name, and a Mac address.</u></p> <p><u>See analysis and evidence for [1pre], limitations [1a] and [1f], and claim 6 above.</u></p>

<u>Claim 31</u>	<u>Accused Instrumentalities</u>
<p><u>[31b] wherein, the plurality of files within each of the plurality of containers comprise one or more application programs including one or more processes, and associated system files for use in executing the one or more processes wherein the associated system files are files that are copies of files or modified copies of files that remain as part of the operating system, each container having its own execution file associated therewith for starting one or more applications, in operation, each container utilizing a kernel resident on the server and wherein each container exclusively uses a kernel in an underlying operation system in which it is running and is absent its own kernel; and,</u></p>	<p><u>Each Accused Instrumentality comprises a system wherein the plurality of files within each of the plurality of containers comprise one or more application programs including one or more processes, and associated system files for use in executing the one or more processes wherein the associated system files are files that are copies of files or modified copies of files that remain as part of the operating system, each container having its own execution file associated therewith for starting one or more applications, in operation, each container utilizing a kernel resident on the server and wherein each container exclusively uses a kernel in an underlying operation system in which it is running and is absent its own kernel.</u></p> <p><u>See analysis and evidence for [1pre], limitations [1a], [1c], [1d], [1e], and [1f], and claim 2 above.</u></p>

~~[31c] a run-time module for monitoring system calls from applications associated with one or more containers and for providing control of the one or more applications.~~

~~Each Accused Instrumentality comprises a run-time module for monitoring system calls from applications associated with one or more containers and for providing control of the one or more applications.~~

~~For example, IBM Cloud Kubernetes Service includes the containerd runtime module or another container runtime. For another example, Kubernetes uses the Linux kernel's seccomp mode to monitor and control system calls made from a container.~~

~~See, e.g.:~~

Security Bulletin: IBM Cloud Kubernetes Service is affected by a containerd security vulnerability (CVE-2024-21626)

Security Bulletin

Summary

IBM Cloud Kubernetes Service is affected by a security vulnerability found in the runc component shipped with containerd where an attacker could gain unauthorized access to the host filesystem (CVE-2024-21626).

~~<https://www.ibm.com/support/pages/security-bulletin-ibm-cloud-kubernetes-service-affected-containerd-security-vulnerability-cve-2024-21626>~~

containerd Adopters

A non-exhaustive list of containerd adopters is provided below.

Docker/Moby engine - Containerd began life prior to its CNCF adoption as a lower-layer runtime manager for `runc` processes below the Docker engine. Continuing today, containerd has extremely broad production usage as a component of the [Docker engine](#) stack. Note that this includes any use of the open source [Moby engine project](#); including the Balena project listed below.

IBM Cloud Kubernetes Service (IKS) - offers containerd as the CRI runtime for v1.11 and higher versions.

IBM Cloud Private (ICP) - IBM's on-premises cloud offering has containerd as a "tech preview" CRI runtime for the Kubernetes offered within this product for the past two releases, and plans to fully migrate to containerd in a future release.

<https://github.com/moby/containerd/blob/docker/20.10/ADOPTERS.md>

Container Runtimes

Note: Dockershim has been removed from the Kubernetes project as of release 1.24. Read the [Dockershim Removal FAQ](#) for further details.

~~You need to install a container runtime into each node in the cluster so that Pods can run there. This page outlines what is involved and describes related tasks for setting up nodes.~~

Kubernetes 1.30 requires that you use a runtime that conforms with the Container Runtime Interface (CRI).

See [CRI version support](#) for more information.

~~<https://kubernetes.io/docs/setup/production-environment/container-runtimes/>~~

Claim 31Accused Instrumentalities

Restrict a Container's Syscalls with seccomp

① **FEATURE STATE:** Kubernetes v1.19 [stable]

Seccomp stands for secure computing mode and has been a feature of the Linux kernel since version 2.6.12. It can be used to sandbox the privileges of a process, restricting the calls it is able to make from userspace into the kernel. Kubernetes lets you automatically apply seccomp profiles loaded onto a node to your Pods and containers.

Identifying the privileges required for your workloads can be difficult. In this tutorial, you will go through how to load seccomp profiles into a local Kubernetes cluster, how to apply them to a Pod, and how you can begin to craft profiles that give only the necessary privileges to your container processes.

<https://kubernetes.io/docs/tutorials/security/seccomp/>

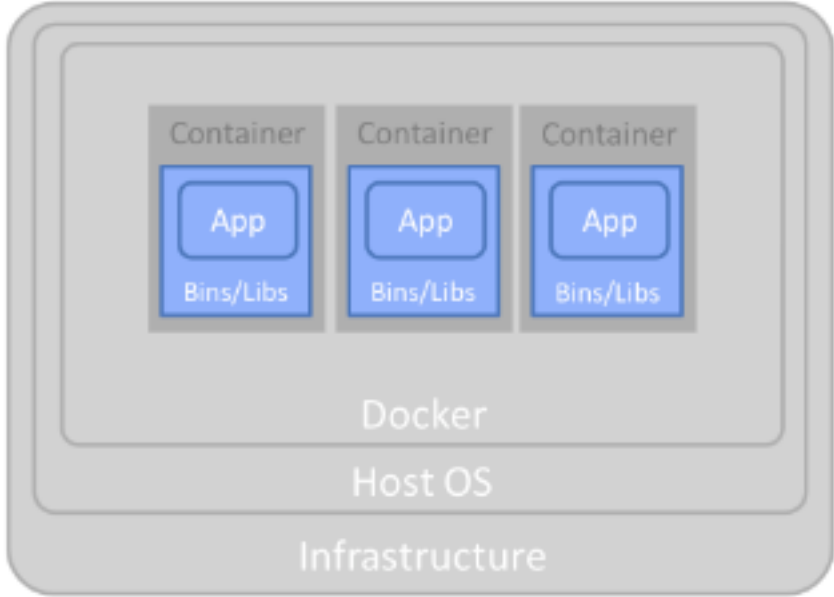
U.S. Patent No. 7,784,058 (“’058 Patent”)

Accused Instrumentalities: IBM products and services using user mode critical system elements as shared libraries, including without limitation IBM Cloud Kubernetes Service (IKS), IBM Cloud Private (ICP), and IBM Hybrid Cloud mesh,, and all versions and variations thereof since the issuance of the asserted patent.

Each Accused Instrumentality infringes the claims in substantially the same way, and the evidence shown in this chart is similarly applicable to each Accused Instrumentality. Each claim limitation is literally infringed by each Accused Instrumentality. However, to the extent any claim limitation is not met literally, it is nonetheless met under the doctrine of equivalents because the differences between the claim limitation and each Accused Instrumentality would be insubstantial, and each Accused Instrumentality performs substantially the same function, in substantially the same way, to achieve the same result as the claimed invention. Notably, Defendant has not yet articulated which, if any, particular claim limitations it believes are not met by the Accused Instrumentalities.

Claim 1






Claim 1	Accused Instrumentalities
[1pre] 1. A computing system for executing a plurality of software applications comprising:	<p>To the extent the preamble is limiting, each Accused Instrumentality comprises or constitutes a computing system for executing a plurality of software applications as claimed.</p> <p><i>See claim limitations below.</i></p> <p><i>See also, e.g.:</i></p> <p>IBM Cloud® Kubernetes Service provides a fully managed container service for Docker (OCI) containers, so clients can deploy containerized apps onto a pool of compute hosts and subsequently manage those containers. Containers are automatically scheduled and placed onto available compute hosts based on your requirements and availability in the cluster.</p> <p>https://www.ibm.com/products/kubernetes-service</p> <p>With IBM Cloud Kubernetes Service, you can deploy Docker containers into pods that run on your worker nodes. The worker nodes come with a set of add-on pods to help you manage your containers. Install more add-ons through Helm, a Kubernetes package manager. These add-ons can extend your apps with dashboards, logging, IBM Cloud and IBM Watson® services and more.</p> <p>https://www.ibm.com/products/kubernetes-service</p>

Claim 1	Accused Instrumentalities
	<p style="text-align: center;">Containers</p>  <p>The diagram illustrates a container architecture stack. At the base is a light gray rounded rectangle labeled 'Infrastructure'. Above it is a slightly darker gray rounded rectangle labeled 'Host OS'. Inside the Host OS is a medium gray rounded rectangle labeled 'Docker'. Within the Docker container, there are three separate gray rounded rectangles, each labeled 'Container'. Each 'Container' contains a blue rounded rectangle labeled 'App' and a smaller blue rounded rectangle below it labeled 'Bins/Libs'.</p> <p style="text-align: center;">https://developer.ibm.com/articles/true-benefits-of-moving-to-containers-1/</p>
[1a] a) a processor;	<p>Each Accused Instrumentality comprises a processor.</p> <p><i>See, e.g.:</i></p>

Claim 1	Accused Instrumentalities
	<p>Containers are executable units of software in which application code is packaged along with its libraries and dependencies, in common ways so that the code can be run anywhere—whether it be on desktop, traditional IT or the cloud.</p> <p>To do this, containers take advantage of a form of operating system (OS) virtualization in which features of the OS kernel (e.g. Linux namespaces and cgroups, Windows silos and job objects) can be leveraged to isolate processes and control the amount of CPU, memory and disk that those processes can access.</p> <p>Containers are small, fast and portable because unlike a virtual machine, containers do not need to include a guest OS in every instance and can instead simply leverage the features and resources of the host OS.</p> <p>https://www.ibm.com/topics/containers</p> <p>Containers use a form of operating system (OS) virtualization. Put simply, they leverage features of the host operating system to isolate processes and control the processes' access to CPUs, memory and desk space.</p> <p>https://www.ibm.com/blog/containers-vs-vms/</p>

Claim 1	Accused Instrumentalities
<p>[1b] b) an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor; and,</p>	<p>Each Accused Instrumentality comprises an operating system having an operating system kernel having OS critical system elements (OSCSEs) for running in kernel mode using said processor.</p> <p><i>See, e.g.:</i></p> <p>Containerization is the packaging of software code with just the operating system (OS) libraries and dependencies required to run the code to create a single lightweight executable—called a container—that runs consistently on any infrastructure. More portable and resource-efficient than virtual machines (VMs), containers have become the de facto compute units of modern cloud-native applications.</p> <p>Containerization allows developers to create and deploy applications faster and more securely. With traditional methods, code is developed in a specific computing environment which, when transferred to a new location, often results in bugs and errors. For example, when a developer transfers code from a desktop computer to a VM or from a Linux to a Windows operating system. Containerization eliminates this problem by bundling the application code together with the related configuration files, libraries, and dependencies required for it to run. This single package of software or “container” is abstracted away from the host operating system, and hence, it stands alone and becomes portable—able to run across any platform or cloud, free of issues.</p> <p>https://www.ibm.com/topics/containerization</p> <p>Kernel mode</p> <p>Kernel mode refers to the processor mode that enables software to have full and unrestricted access to the system and its resources. The OS kernel and kernel drivers, such as the file system driver, are loaded into protected memory space and operate in this highly privileged kernel mode.</p> <p>https://www.techtarget.com/searchdatacenter/definition/kernel</p>

Claim 1	Accused Instrumentalities
	<p>The GNU C Library, commonly known as glibc, is the GNU Project implementation of the C standard library. It is a wrapper around the system calls of the Linux kernel for application use. Despite its name, it now also directly supports C++ (and, indirectly, other programming languages). It was started in the 1980s by the Free Software Foundation (FSF) for the GNU operating system.</p> <p>https://en.wikipedia.org/wiki/Glibc</p>
<p>[1c] c) a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode and</p>	<p>Each Accused Instrumentality comprises a shared library having shared library critical system elements (SLCSEs) stored therein for use by the plurality of software applications in user mode.</p> <p><i>See, e.g.:</i></p> <p>A Docker image is the basis for every container that you create with IBM Cloud® Kubernetes Service.</p> <p>An image is created from a Dockerfile, which is a file that contains instructions to build the image. A Dockerfile might reference build artifacts in its instructions that are stored separately, such as an app, the app's configuration, and its dependencies.</p> <p>https://cloud.ibm.com/docs/containers?topic=containers-images</p>

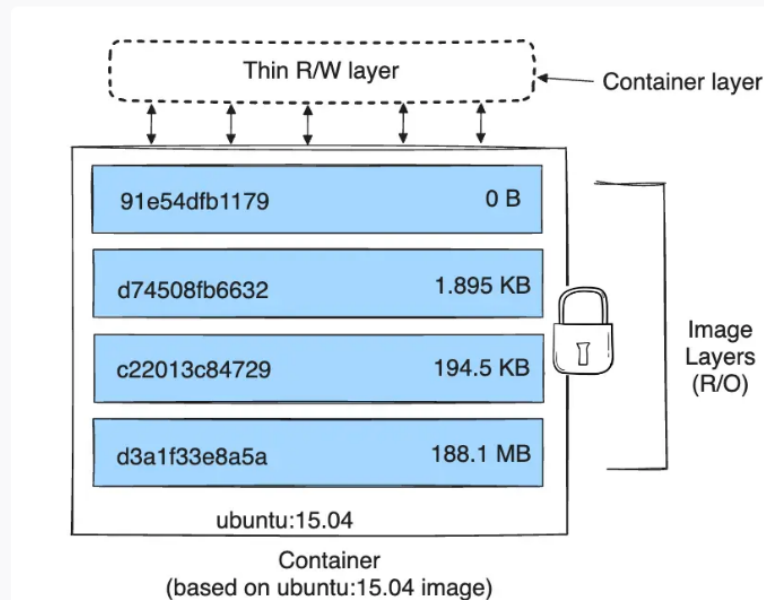
Claim 1	Accused Instrumentalities		
	Docker base image	Supported versions	Source of security notices
	Alpine	All stable versions with vendor security support.	Alpine SecDB database  .
	Debian	All stable versions with vendor security support. CVEs on binary packages that are associated with the Debian source package <code>linux</code> , such as <code>linux-libc-dev</code> , are not reported. Most of these binary packages are kernel and kernel modules, which are not run in container images.	Debian Security Bug Tracker  .
	GoogleContainerTools distroless	All stable versions with vendor security support.	GoogleContainerTools distroless  .
	Red Hat® Enterprise Linux® (RHEL)	RHEL/UBI 7, RHEL/UBI 8, and RHEL/UBI 9	Red Hat Security Data API  .
	Ubuntu	All stable versions with vendor security support.	Ubuntu CVE Tracker  .
	https://cloud.ibm.com/docs/Registry?topic=Registry-va_index&interface=ui		

Claim 1	Accused Instrumentalities
	<h2 data-bbox="646 280 1272 349">About storage drivers</h2> <p data-bbox="646 394 1871 521">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="646 589 1564 646">Storage drivers versus Docker volumes</h2> <p data-bbox="646 683 1913 948">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="646 1000 1902 1122">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1339 1049 1535 1073" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="636 1154 1226 1187"><a data-bbox="636 1154 1226 1187" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="657 280 1081 334">Images and layers</h2> <p data-bbox="657 371 1822 448">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="674 516 1453 829"> # syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="657 894 1900 1203">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="634 1224 1226 1256">https://docs.docker.com/storage/storagedriver/</p>

Each layer is only a set of differences from the layer before it. Note that both *adding*, and *removing* files will result in a new layer. In the example above, the `$HOME/.cache` directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the [Best practices for writing Dockerfiles](#) and [use multi-stage builds](#) sections to learn how to optimize your Dockerfiles for efficient images.

The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an `ubuntu:15.04` image.






<https://docs.docker.com/storage/storagedriver/>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 289 919 354">Volumes</h2> <p data-bbox="653 407 1906 537">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="634 558 1308 591">https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2 data-bbox="653 643 1226 691">Container environment</h2> <p data-bbox="653 729 1474 797">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="695 834 1451 992" style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster. <p data-bbox="634 1029 1528 1062">https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="659 282 877 345">Images</h2> <p data-bbox="659 378 1522 532">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="659 568 1528 641">You typically create a container image of your application and push it to a registry before referring to it in a <u>Pod</u>.</p> <p data-bbox="634 667 1329 699">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="653 743 919 800">Volumes</h2> <p data-bbox="653 841 1482 914">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers.</p> <p data-bbox="653 922 1430 954">One problem occurs when a container crashes or is stopped.</p> <p data-bbox="653 963 1528 1287">Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a <code>Pod</code> and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes <u>volume</u> abstraction solves both of these problems. Familiarity with <u>Pods</u> is suggested.</p> <p data-bbox="634 1312 1308 1344">https://kubernetes.io/docs/concepts/storage/volumes/</p>

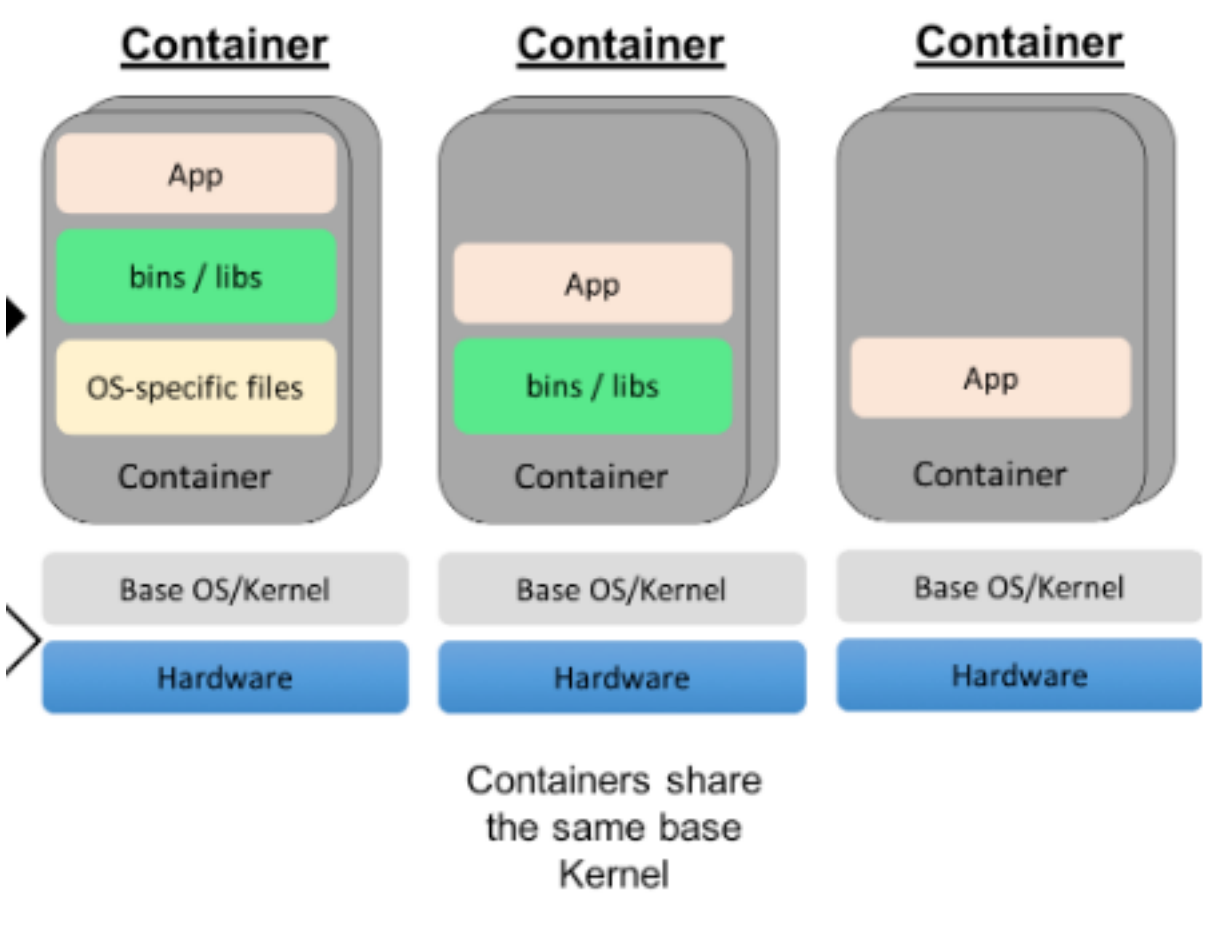
Claim 1	Accused Instrumentalities
	<h2 data-bbox="674 289 1297 347">Open Container Initiative</h2> <hr data-bbox="674 358 1911 362"/> <h3 data-bbox="674 412 1182 454">Image Format Specification</h3> <hr data-bbox="674 467 1911 470"/> <p data-bbox="674 505 1892 578">This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p> <p data-bbox="674 613 1902 686">The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p> <p data-bbox="634 716 1478 784">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<div><h3>Overview</h3><p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p><div><div><pre>public class HelloWorld { public static void main(String[] args) { System.out.println("Hello, World"); } }</pre></div><div>→</div><div><div><div>/bin/java /opt/app.jar /lib/libc</div><div>layer</div></div><div>+</div><div><div><div>{ "manifests": { "platform": { "os": "linux", ... } } }</div><div>image index</div></div><div>+</div><div><div><div>{ ... "config": { "Cmd": ["java", "-jar", "app.jar"], ... } }</div><div>config</div></div></div></div><div>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</div></div></div></div>

Claim 1	Accused Instrumentalities
	<div data-bbox="653 272 1297 337"><h2>OCI Image Configuration</h2></div> <div data-bbox="653 386 1915 553"><p>An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p></div> <div data-bbox="653 586 1661 625"><p>This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p></div> <div data-bbox="634 651 1505 727"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p></div>



Claim 1	Accused Instrumentalities
	<p data-bbox="661 284 745 316">Layer</p> <ul data-bbox="688 357 1921 706" style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p data-bbox="661 755 856 787">Image JSON</p> <ul data-bbox="688 828 1921 1177" style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p data-bbox="634 1201 1501 1274">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
	<p><i>Docker images</i> contain executable application source code as well as all the tools, libraries, and dependencies that the application code needs to run as a container. When you run the Docker image, it becomes one instance (or multiple instances) of the container.</p> <p>It's possible to build a Docker image from scratch, but most developers pull them down from common repositories. Multiple Docker images can be created from a single base image, and they'll share the commonalities of their stack.</p> <p>Docker images are made up of layers, and each layer corresponds to a version of the image. Whenever a developer makes changes to the image, a new top layer is created, and this top layer replaces the previous top layer as the current version of the image. Previous layers are saved for rollbacks or to be re-used in other projects.</p> <p>Each time a container is created from a Docker image, yet another new layer called the container layer is created. Changes made to the container—such as the addition or deletion of files—are saved to the container layer only and exist only while the container is running. This iterative image-creation process enables increased overall efficiency since multiple live container instances can run from just a single base image, and when they do so, they leverage a common stack.</p> <p>https://www.ibm.com/topics/docker</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="630 292 1827 1218"><p><u>Container</u> <u>Container</u> <u>Container</u></p><p>The diagram illustrates three containers, each represented by a rounded rectangle. The first container on the left contains three stacked boxes: 'App' (orange), 'bins / libs' (green), and 'OS-specific files' (yellow). The second container in the middle contains two stacked boxes: 'App' (orange) and 'bins / libs' (green). The third container on the right contains one box: 'App' (orange). Below each container is a grey box labeled 'Base OS/Kernel'. At the bottom of the stack is a blue box labeled 'Hardware'. A large bracket on the left side of the 'Base OS/Kernel' and 'Hardware' boxes indicates they are shared across all three containers. Below the diagram, the text 'Containers share the same base Kernel' is centered.</p><p>Containers share the same base Kernel</p></div> <div data-bbox="630 1266 1024 1307">https://ibm.github.io/kube101/</div>

Claim 1	Accused Instrumentalities
	<p>The GNU C Library, commonly known as glibc, is the GNU Project implementation of the C standard library. It is a wrapper around the system calls of the Linux kernel for application use. Despite its name, it now also directly supports C++ (and, indirectly, other programming languages). It was started in the 1980s by the Free Software Foundation (FSF) for the GNU operating system.</p> <p>https://en.wikipedia.org/wiki/Glibc</p>
<p>[1d] i) wherein some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible to some of the plurality of software applications and when one of the SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications,</p>	<p>In each Accused Instrumentality, some of the SLCSEs stored in the shared library are functional replicas of OSCSEs and are accessible to some of the plurality of software applications and when one of the SLCSEs is accessed by one or more of the plurality of software applications it forms a part of the one or more of the plurality of software applications.</p> <p>For example, a base image serves as a self-contained unit that encompasses all the necessary components for an application to run, including the application code, runtime environment, system tools, and dependencies (i.e., SLCSEs). The images are based on existing Linux distributions, such as Debian and Ubuntu, including essential system elements (i.e., functional replicas of OSCSEs). Each container image is based on a specific base image, which contains the application code, and dependencies, including system libraries or shared library critical system elements (SLCSEs). When the container runs the image, it creates a runtime instance of that container image.</p> <p><i>See, e.g.:</i></p> <p>A Docker image is the basis for every container that you create with IBM Cloud® Kubernetes Service.</p> <p>An image is created from a Dockerfile, which is a file that contains instructions to build the image. A Dockerfile might reference build artifacts in its instructions that are stored separately, such as an app, the app's configuration, and its dependencies.</p> <p>https://cloud.ibm.com/docs/containers?topic=containers-images</p>

Claim 1	Accused Instrumentalities
	<p><i>Docker images</i> contain executable application source code as well as all the tools, libraries, and dependencies that the application code needs to run as a container. When you run the Docker image, it becomes one instance (or multiple instances) of the container.</p> <p>It's possible to build a Docker image from scratch, but most developers pull them down from common repositories. Multiple Docker images can be created from a single base image, and they'll share the commonalities of their stack.</p> <p>Docker images are made up of layers, and each layer corresponds to a version of the image. Whenever a developer makes changes to the image, a new top layer is created, and this top layer replaces the previous top layer as the current version of the image. Previous layers are saved for rollbacks or to be re-used in other projects.</p> <p>Each time a container is created from a Docker image, yet another new layer called the container layer is created. Changes made to the container—such as the addition or deletion of files—are saved to the container layer only and exist only while the container is running. This iterative image-creation process enables increased overall efficiency since multiple live container instances can run from just a single base image, and when they do so, they leverage a common stack.</p> <p>https://www.ibm.com/topics/docker</p> <p>Following software is installed on the Docker containers as part of the Product Master image deployment:</p> <ul style="list-style-type: none"> – Red Hat Enterprise Linux (RHEL) 7 Universal Base Image (UBI) base Docker image <p>https://www.ibm.com/docs/en/product-master/12.0.0?topic=deployment-installing-product-by-using-docker-images</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="646 277 1892 493"><div><div>ubuntu</div><div>Updated 15 days ago</div><div>Ubuntu is a Debian-based Linux operating system based on free software.</div><div>Linux IBM Z 386 riscv64 x86-64 ARM ARM 64 PowerPC 64 LE</div></div><div data-bbox="1692 313 1839 332">1B+ · 10K+</div></div> <div data-bbox="646 540 1892 756"><div><div>debian</div><div>Updated 35 minutes ago</div><div>Debian is a Linux distribution that's composed entirely of free and open-source software.</div><div>Linux riscv64 x86-64 ARM ARM 64 386 mips64le PowerPC 64 LE IBM Z</div></div><div data-bbox="1709 579 1852 599">1B+ · 4.9K</div></div> <div data-bbox="634 803 1533 836">https://hub.docker.com/search?image_filter=official&type=image&q=</div>

Claim 1	Accused Instrumentalities																																																						
	<table><tr><th>Platform</th><th>x86_64 / amd64</th><th>arm64 / aarch64</th><th>arm (32-bit)</th><th>ppc64le</th><th>s390x</th></tr><tr><td>CentOS</td><td>✓</td><td>✓</td><td></td><td>✓</td><td></td></tr><tr><td>Debian</td><td>✓</td><td>✓</td><td>✓</td><td>✓</td><td></td></tr><tr><td>Fedora</td><td>✓</td><td>✓</td><td></td><td>✓</td><td></td></tr><tr><td>Raspberry Pi OS (32-bit)</td><td></td><td></td><td>✓</td><td></td><td></td></tr><tr><td>RHEL (s390x)</td><td></td><td></td><td></td><td></td><td>✓</td></tr><tr><td>SLES</td><td></td><td></td><td></td><td></td><td>✓</td></tr><tr><td>Ubuntu</td><td>✓</td><td>✓</td><td>✓</td><td>✓</td><td>✓</td></tr><tr><td>Binaries</td><td>✓</td><td>✓</td><td>✓</td><td></td><td></td></tr></table> <p>https://docs.docker.com/engine/install/</p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p>https://www.techtarget.com/searchitoperations/definition/Docker-image</p> <p><i>Docker images</i> contain executable application source code as well as all the tools, libraries, and dependencies that the application code needs to run as a container. When you run the Docker image, it becomes one instance (or multiple instances) of the container.</p> <p>https://www.ibm.com/topics/docker</p>	Platform	x86_64 / amd64	arm64 / aarch64	arm (32-bit)	ppc64le	s390x	CentOS	✓	✓		✓		Debian	✓	✓	✓	✓		Fedora	✓	✓		✓		Raspberry Pi OS (32-bit)			✓			RHEL (s390x)					✓	SLES					✓	Ubuntu	✓	✓	✓	✓	✓	Binaries	✓	✓	✓		
Platform	x86_64 / amd64	arm64 / aarch64	arm (32-bit)	ppc64le	s390x																																																		
CentOS	✓	✓		✓																																																			
Debian	✓	✓	✓	✓																																																			
Fedora	✓	✓		✓																																																			
Raspberry Pi OS (32-bit)			✓																																																				
RHEL (s390x)					✓																																																		
SLES					✓																																																		
Ubuntu	✓	✓	✓	✓	✓																																																		
Binaries	✓	✓	✓																																																				

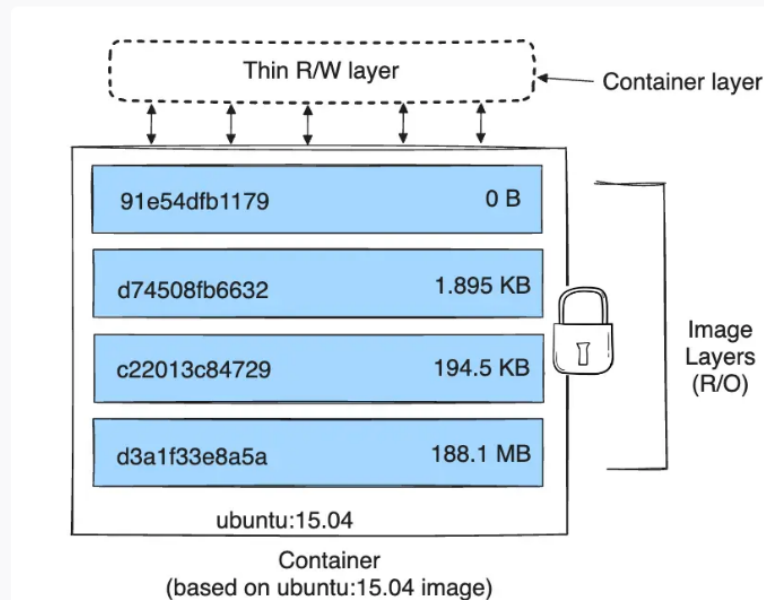
Claim 1	Accused Instrumentalities
	<p data-bbox="892 293 1587 329">A container is a runtime instance of a docker image.</p> <p data-bbox="892 386 1304 418">A Docker container consists of</p> <div data-bbox="659 475 1350 651"><p data-bbox="659 529 789 561">container</p><ul style="list-style-type: none"><li data-bbox="909 475 1163 508">• A Docker image<li data-bbox="909 548 1304 581">• An execution environment<li data-bbox="909 621 1350 651">• A standard set of instructions</div> <p data-bbox="636 711 1157 747">https://docs.docker.com/glossary/#image</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="646 280 1272 349">About storage drivers</h2> <p data-bbox="646 394 1871 521">To use storage drivers effectively, it's important to know how Docker builds and stores images, and how these images are used by containers. You can use this information to make informed choices about the best way to persist data from your applications and avoid performance problems along the way.</p> <h2 data-bbox="646 586 1564 646">Storage drivers versus Docker volumes</h2> <p data-bbox="646 683 1913 950">Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p data-bbox="646 998 1902 1125">Use Docker volumes for write-intensive data, data that must persist beyond the container's lifespan, and data that must be shared between containers. Refer to the <a data-bbox="1339 1047 1535 1076" href="#">volumes section to learn how to use volumes to persist data and improve performance.</p> <p data-bbox="636 1154 1224 1187"><a data-bbox="636 1154 1224 1187" href="https://docs.docker.com/storage/storagedriver/">https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 277 1083 334">Images and layers</h2> <p data-bbox="653 370 1822 448">A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre data-bbox="653 488 1908 854"> # syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py </pre> <p data-bbox="653 894 1900 1203">This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p> <p data-bbox="632 1224 1226 1256">https://docs.docker.com/storage/storagedriver/</p>

Each layer is only a set of differences from the layer before it. Note that both *adding*, and *removing* files will result in a new layer. In the example above, the `$HOME/.cache` directory is removed, but will still be available in the previous layer and add up to the image's total size. Refer to the [Best practices for writing Dockerfiles](#) and [use multi-stage builds](#) sections to learn how to optimize your Dockerfiles for efficient images.

The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an `ubuntu:15.04` image.

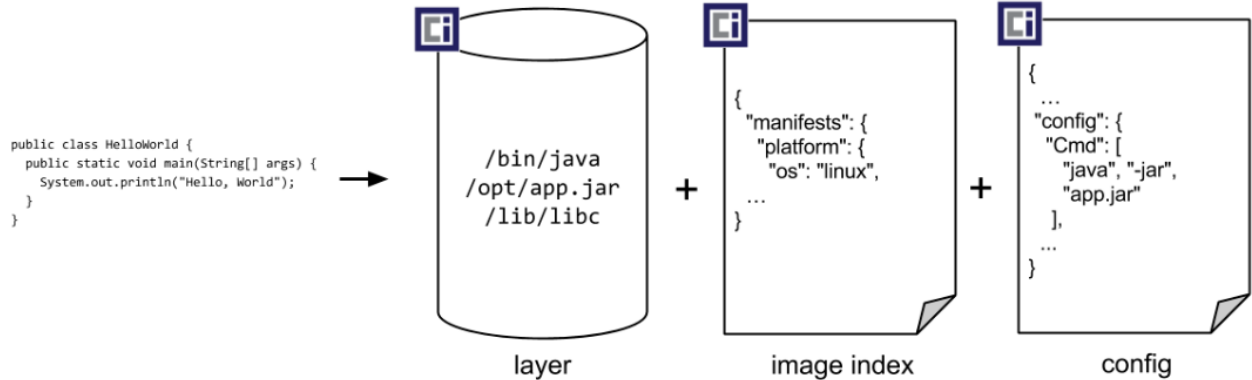


<https://docs.docker.com/storage/storagedriver/>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 289 919 354">Volumes</h2> <p data-bbox="653 407 1906 537">Volumes are the preferred mechanism for persisting data generated by and used by Docker containers. While bind mounts are dependent on the directory structure and OS of the host machine, volumes are completely managed by Docker. Volumes have several advantages over bind mounts:</p> <p data-bbox="636 558 1308 591">https://kubernetes.io/docs/concepts/storage/volumes/</p> <h2 data-bbox="653 643 1226 691">Container environment</h2> <p data-bbox="653 729 1474 797">The Kubernetes Container environment provides several important resources to Containers:</p> <ul data-bbox="695 834 1451 992" style="list-style-type: none">• A filesystem, which is a combination of an image and one or more volumes.• Information about the Container itself.• Information about other objects in the cluster. <p data-bbox="636 1029 1528 1062">https://kubernetes.io/docs/concepts/containers/container-environment/</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="659 285 877 347">Images</h2> <p data-bbox="659 380 1522 532">A container image represents binary data that encapsulates an application and all its software dependencies. Container images are executable software bundles that can run standalone and that make very well defined assumptions about their runtime environment.</p> <p data-bbox="659 570 1528 641">You typically create a container image of your application and push it to a registry before referring to it in a <u>Pod</u>.</p> <p data-bbox="634 667 1329 703">https://kubernetes.io/docs/concepts/containers/images/</p> <h2 data-bbox="653 743 919 805">Volumes</h2> <p data-bbox="653 841 1482 912">On-disk files in a container are ephemeral, which presents some problems for non-trivial applications when running in containers.</p> <p data-bbox="653 922 1430 954">One problem occurs when a container crashes or is stopped.</p> <p data-bbox="653 964 1528 1284">Container state is not saved so all of the files that were created or modified during the lifetime of the container are lost. During a crash, kubelet restarts the container with a clean state. Another problem occurs when multiple containers are running in a <code>Pod</code> and need to share files. It can be challenging to setup and access a shared filesystem across all of the containers. The Kubernetes <u>volume</u> abstraction solves both of these problems. Familiarity with <code>Pods</code> is suggested.</p> <p data-bbox="634 1310 1308 1346">https://kubernetes.io/docs/concepts/storage/volumes/</p>

Claim 1	Accused Instrumentalities
	<div data-bbox="669 289 1297 347"><h2>Open Container Initiative</h2><hr/></div> <div data-bbox="669 410 1184 456"><h3>Image Format Specification</h3><hr/></div> <div data-bbox="669 503 1890 576"><p>This specification defines an OCI Image, consisting of an image manifest, an image index (optional), a set of filesystem layers, and a configuration.</p></div> <div data-bbox="669 613 1900 686"><p>The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.</p></div> <div data-bbox="632 714 1478 784"><p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p></div>

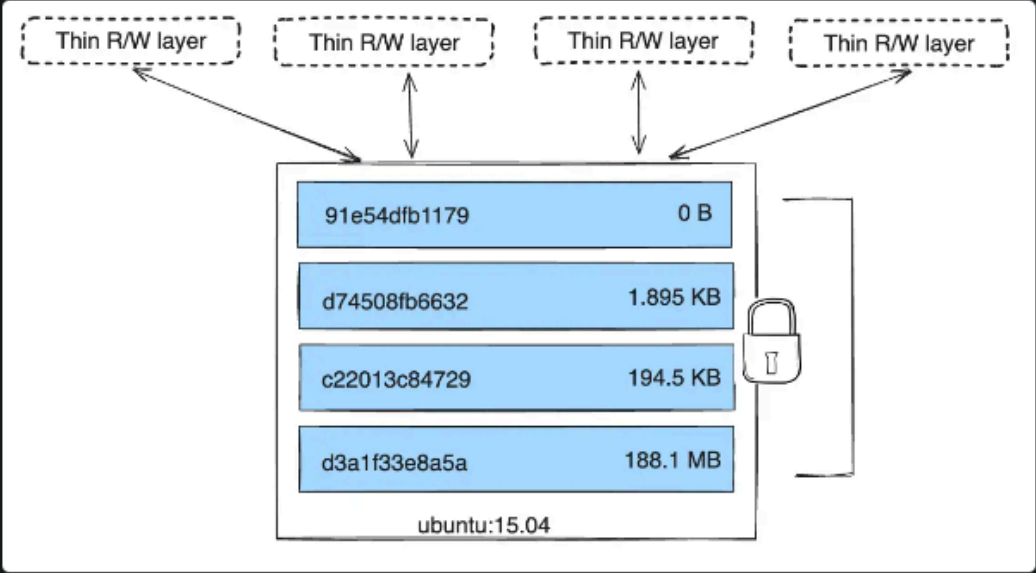
Claim 1	Accused Instrumentalities
	<p>Overview</p> <p>At a high level the image manifest contains metadata about the contents and dependencies of the image including the content-addressable identity of one or more filesystem layer changeset archives that will be unpacked to make up the final runnable filesystem. The image configuration includes information such as application arguments, environments, etc. The image index is a higher-level manifest which points to a list of manifests and descriptors. Typically, these manifests may provide different implementations of the image, possibly varying by platform or other attributes.</p>  <p>https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/spec.md</p>

Claim 1	Accused Instrumentalities
	<h2 data-bbox="653 277 1297 337">OCI Image Configuration</h2> <p data-bbox="653 391 1913 553">An OCI <i>Image</i> is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. This specification outlines the JSON format describing images for use with a container runtime and execution tool and its relationship to filesystem changesets, described in Layers.</p> <p data-bbox="653 591 1661 623">This section defines the <code>application/vnd.oci.image.config.v1+json</code> media type.</p> <p data-bbox="632 656 1507 727">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

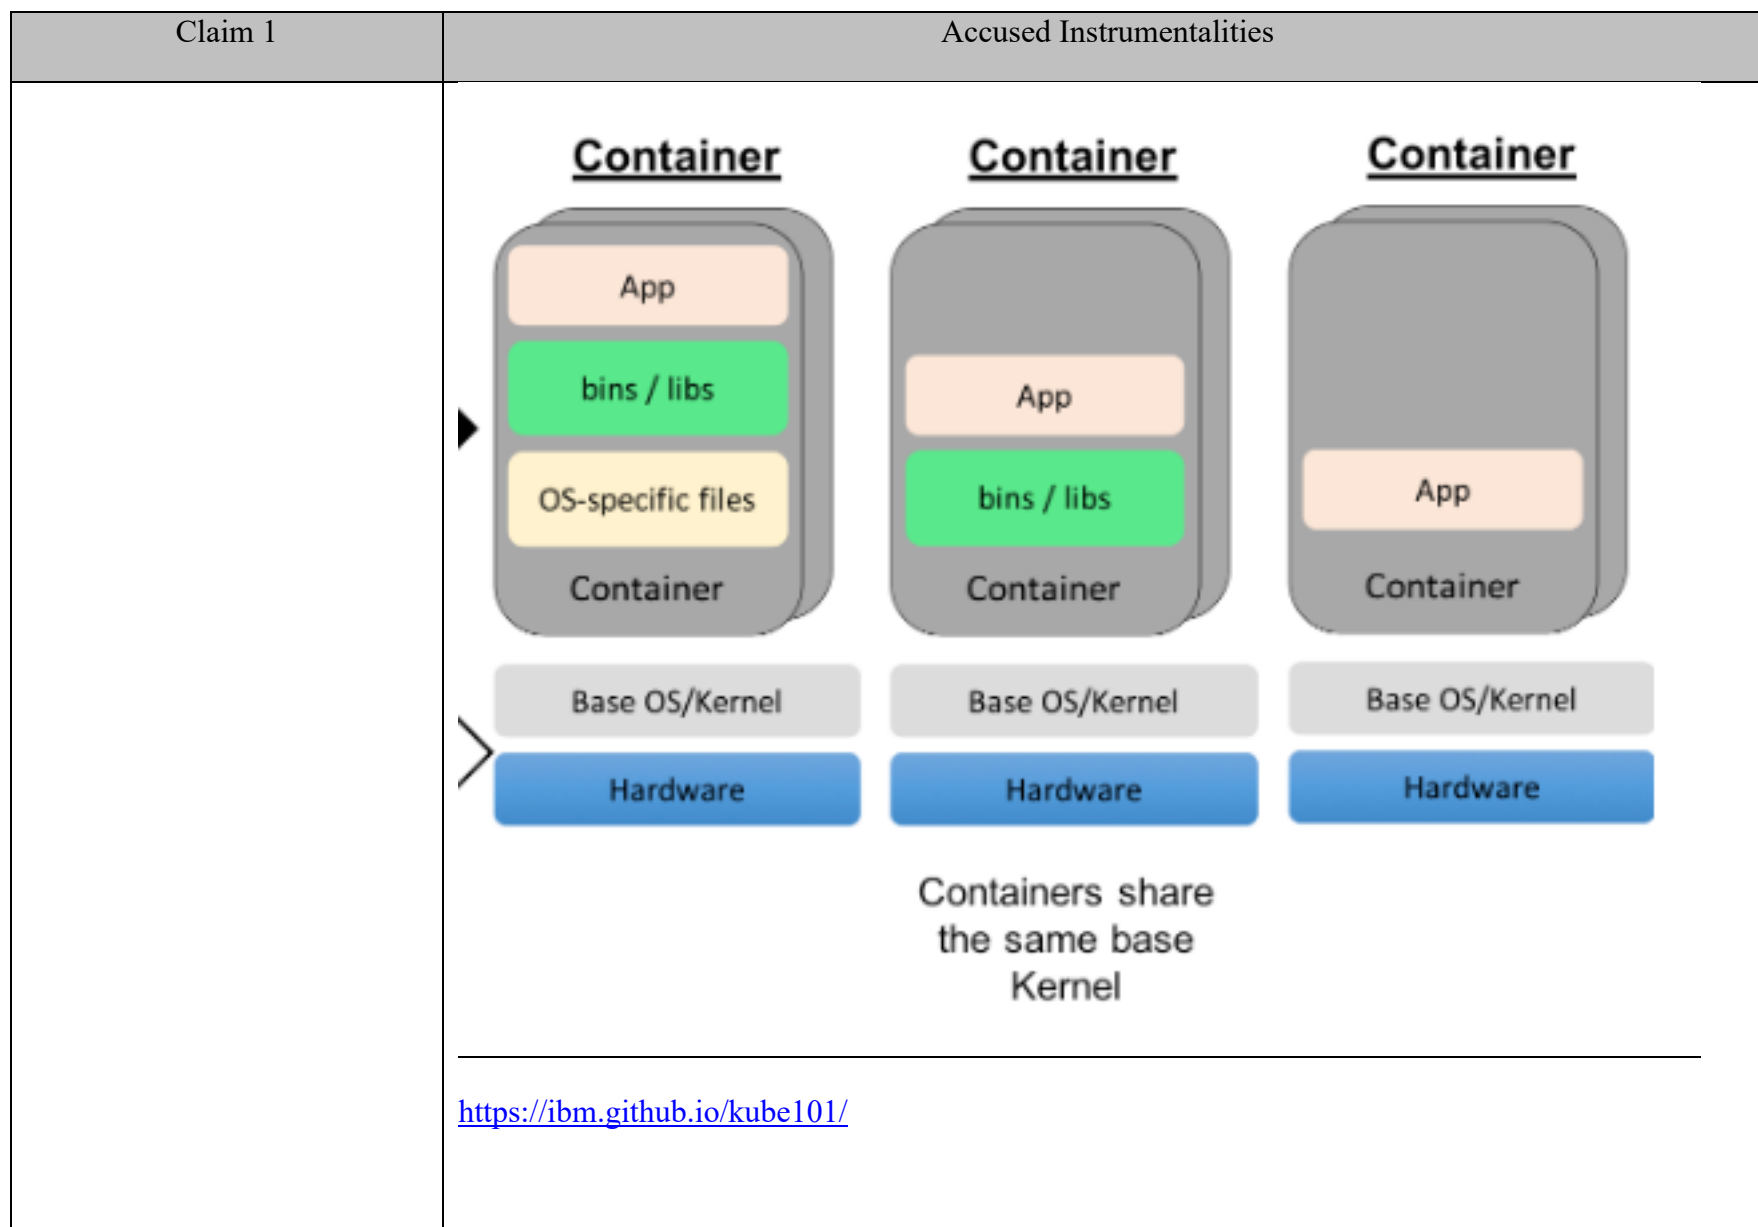
Claim 1	Accused Instrumentalities
	<p data-bbox="661 284 745 316">Layer</p> <ul data-bbox="688 357 1921 706" style="list-style-type: none"> • Image filesystems are composed of <i>layers</i>. • Each layer represents a set of filesystem changes in a tar-based layer format, recording files to be added, changed, or deleted relative to its parent layer. • Layers do not have configuration metadata such as environment variables or default arguments - these are properties of the image as a whole rather than any particular layer. • Using a layer-based or union filesystem such as AUFS, or by computing the diff from filesystem snapshots, the filesystem changeset can be used to present a series of image layers as if they were one cohesive filesystem. <p data-bbox="661 755 856 787">Image JSON</p> <ul data-bbox="688 828 1921 1177" style="list-style-type: none"> • Each image has an associated JSON structure which describes some basic information about the image such as date created, author, as well as execution/runtime configuration like its entrypoint, default arguments, networking, and volumes. • The JSON structure also references a cryptographic hash of each layer used by the image, and provides history information for those layers. • This JSON is considered to be immutable, because changing it would change the computed ImageID. • Changing it means creating a new derived image, instead of changing the existing image. <p data-bbox="634 1201 1501 1274">https://github.com/opencontainers/image-spec/blob/a6af2b480dcfc001ba975f44de53001c873cb0ef/config.md</p>

Claim 1	Accused Instrumentalities
<p>[1e] ii) wherein an instance of a SLCSE provided to at least a first of the plurality of software applications from the shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the operating system have use of a unique instance of a corresponding critical system element for performing same function, and</p>	<p>In each Accused Instrumentality, an instance of a SLCSE provided to at least a first of the plurality of software applications from the shared library is run in a context of said at least first of the plurality of software applications without being shared with other of the plurality of software applications and where at least a second of the plurality of software applications running under the operating system have use of a unique instance of a corresponding critical system element for performing same function.</p> <p>When a Docker or Kubernetes image is used to create a container, it creates a separate and isolated instance of a runtime environment which is independent of other containers running on the same host. Each container has its own instance of base images and its own data. The containers run in isolation, ensuring that the SLCSEs stored in the shared library are accessible to the software applications running in their respective containers. The image includes essential system files, libraries, and dependencies required to run the software application within the container. The containers can share common dependencies and components using layered images. This means that multiple containers utilize the same base image to create an instance. When an instance of SLCSE is provided from the base image (i.e., from the shared library) to an individual container including application software, it operates in isolation and runs its own instance of the software application without sharing resources or critical system elements with other containers. This ensures that each container has its own isolated context. Docker or Kubernetes containers can share common dependencies and components using layered images. This means that multiple containers can utilize the same base image. Therefore, each container, containing the application software running under the operating system, utilizes a unique instance of the corresponding critical system element to execute the respective application software for performing a same or a different function.</p> <p><i>See, e.g.:</i></p>

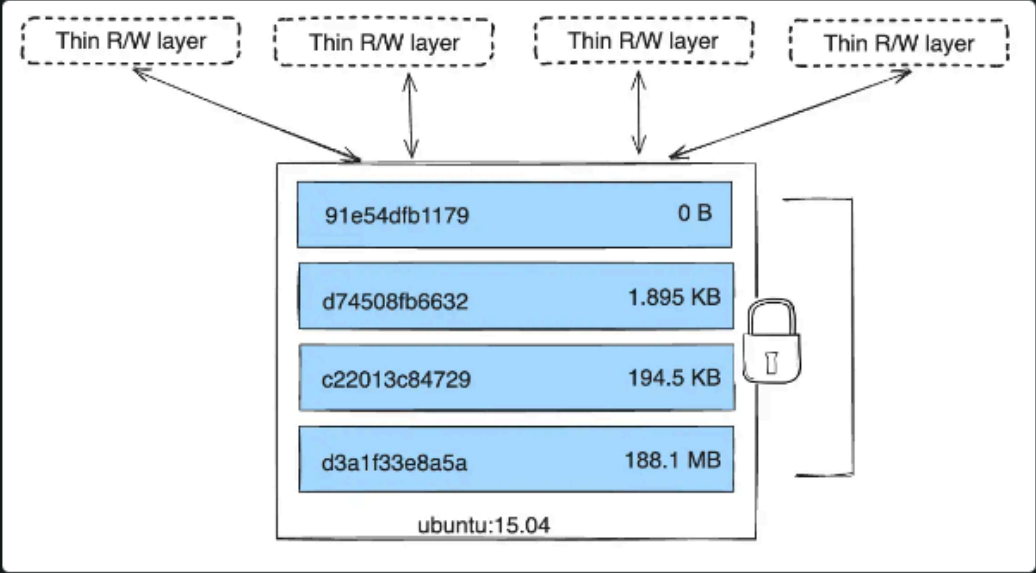
Claim 1	Accused Instrumentalities
	<p><i>Docker images</i> contain executable application source code as well as all the tools, libraries, and dependencies that the application code needs to run as a container. When you run the Docker image, it becomes one instance (or multiple instances) of the container.</p> <p>It's possible to build a Docker image from scratch, but most developers pull them down from common repositories. Multiple Docker images can be created from a single base image, and they'll share the commonalities of their stack.</p> <p>https://www.ibm.com/topics/docker</p>

Claim 1	Accused Instrumentalities
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p>https://docs.docker.com/storage/storagedriver/</p>

Claim 1	Accused Instrumentalities
	<p><i>Docker images</i> contain executable application source code as well as all the tools, libraries, and dependencies that the application code needs to run as a container. When you run the Docker image, it becomes one instance (or multiple instances) of the container.</p> <p>https://www.ibm.com/topics/docker</p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p>https://www.techtarget.com/searchitoperations/definition/Docker-image</p>



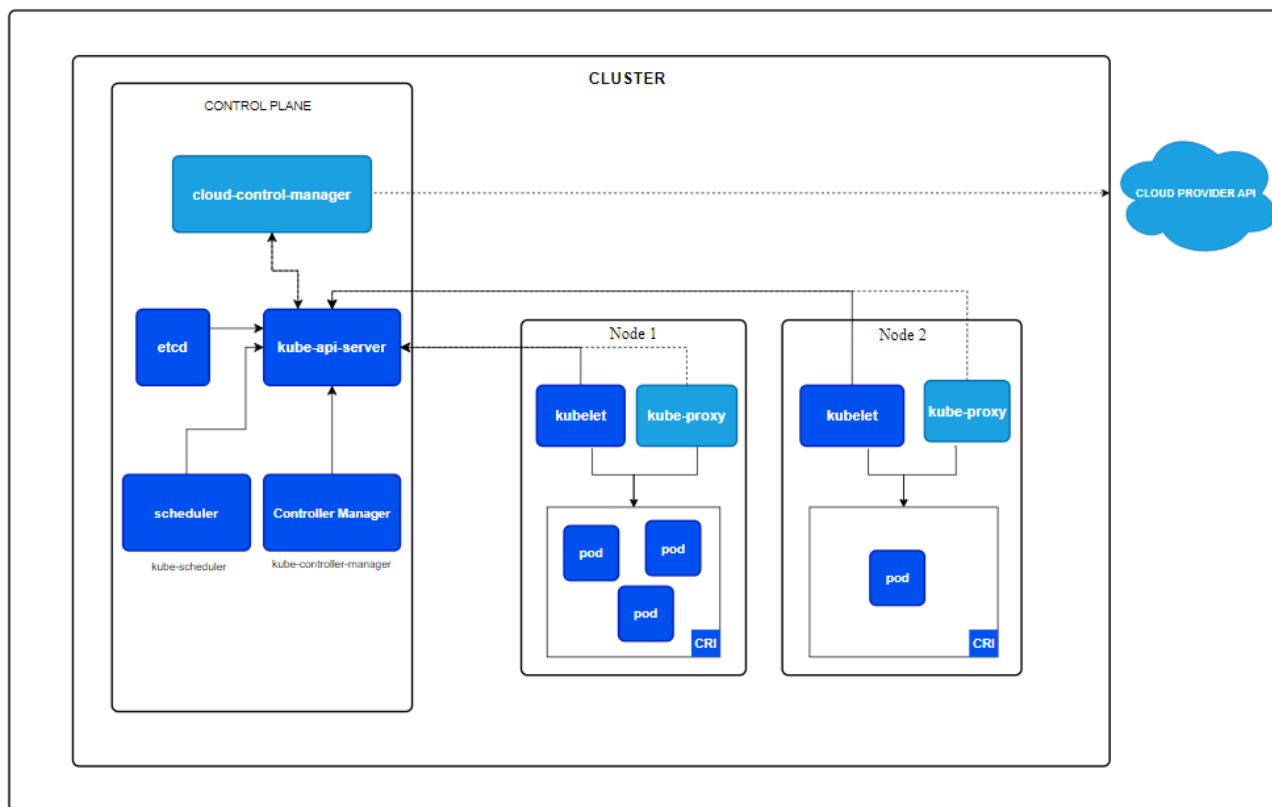
Claim 1	Accused Instrumentalities
<p>[1f] iii) wherein a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.</p>	<p>In each Accused Instrumentality, a SLCSE related to a predetermined function is provided to the first of the plurality of software applications for running a first instance of the SLCSE, and wherein a SLCSE for performing a same function is provided to the second of the plurality of software applications for running a second instance of the SLCSE simultaneously.</p> <p>For example, in Docker or Kubernetes containers, each container operates independently, and a base image includes essential system files, libraries, and dependencies (i.e., SLCSEs) required to run the software application within the container. Based on information and belief, each element, such as system files, libraries, and dependencies (i.e., SLCSE) is associated with an execution of a predetermined function related to the application. When an image is used to create a container in the Accused Instrumentality, an instance of the SLCSE is provided to a software application. Therefore, different instances of the SLCSE are provided to different applications for performing either a same or a different function, simultaneously.</p> <p><i>See, e.g.:</i></p> <p>Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p>https://www.techtarget.com/searchitoperations/definition/Docker-image</p> <p><i>Docker images</i> contain executable application source code as well as all the tools, libraries, and dependencies that the application code needs to run as a container. When you run the Docker image, it becomes one instance (or multiple instances) of the container.</p> <p>https://www.ibm.com/topics/docker</p>

Claim 1	Accused Instrumentalities
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p>https://docs.docker.com/storage/storagedriver/</p> <p>A container is a runnable instance of an image. You can create, start, stop, move, or delete a container using the Docker API or CLI. You can connect a container to one or more networks, attach storage to it, or even create a new image based on its current state.</p>

Claim 1	Accused Instrumentalities
	https://docs.docker.com/get-started/overview/

Claim 2

Claim 2	Accused Instrumentalities
2. A computing system as defined in claim 1, wherein in operation, multiple instances of an SLCSE stored in the shared library run simultaneously within the operating system.	<p>Each Accused Instrumentality comprises or constitutes a computing system as defined in claim 1, wherein in operation, multiple instances of an SLCSE stored in the shared library run simultaneously within the operating system.</p> <p>For example, an individual host/node runs multiple containers and/or pods simultaneously, each of which has an instance of an SLCSE.</p> <p><i>See, e.g.:</i></p>

Claim 2**Accused Instrumentalities**

Kubernetes cluster architecture

<https://kubernetes.io/docs/concepts/architecture/>






Claim 2	Accused Instrumentalities
	<h1 data-bbox="661 256 1138 337">Containers</h1> <p data-bbox="661 396 1915 553">Each container that you run is repeatable; the standardization from having dependencies included means that you get the same behavior wherever you run it.</p> <p data-bbox="661 612 1915 769">Containers decouple applications from the underlying host infrastructure. This makes deployment easier in different cloud or OS environments.</p> <p data-bbox="661 828 1892 985">Each <u>node</u> in a Kubernetes cluster runs the containers that form the Pods assigned to that node. Containers in a Pod are co-located and co-scheduled to run on the same node.</p> <p data-bbox="646 1036 1241 1068">https://kubernetes.io/docs/concepts/containers/</p>

Claim 2	Accused Instrumentalities
	<h1 data-bbox="674 245 1520 318">Kubernetes Scheduler</h1> <p data-bbox="674 370 1650 456">In Kubernetes, <i>scheduling</i> refers to making sure that <u>Pods</u> are matched to <u>Nodes</u> so that <u>Kubelet</u> can run them.</p> <h2 data-bbox="674 578 1325 651">Scheduling overview</h2> <p data-bbox="674 695 1745 935">A scheduler watches for newly created Pods that have no Node assigned. For every Pod that the scheduler discovers, the scheduler becomes responsible for finding the best Node for that Pod to run on. The scheduler reaches this placement decision taking into account the scheduling principles described below.</p> <p data-bbox="674 984 1696 1127">If you want to understand why Pods are placed onto a particular Node, or if you're planning to implement a custom scheduler yourself, this page will help you learn about scheduling.</p> <p data-bbox="642 1175 1566 1208">https://kubernetes.io/docs/concepts/scheduling-eviction/kube-scheduler/</p>

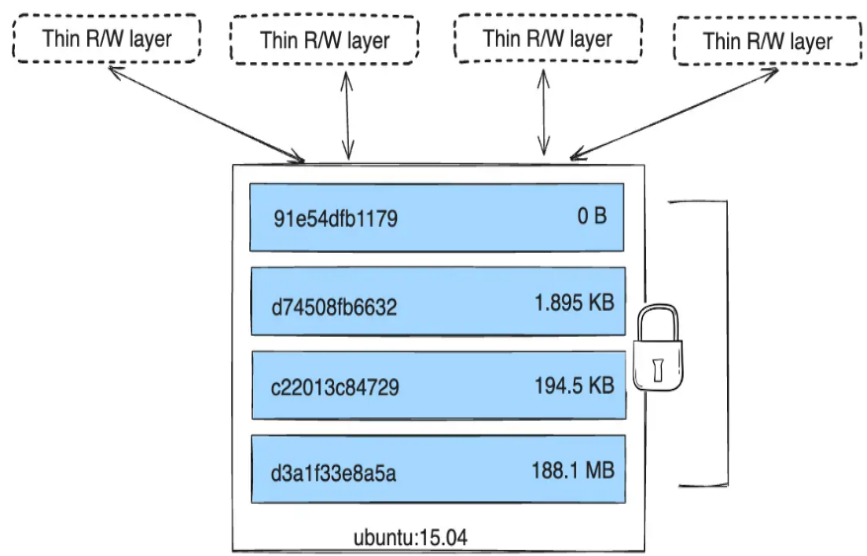
Claim 2	Accused Instrumentalities
	<h2 data-bbox="657 248 1220 313">Running containers</h2> <p data-bbox="657 362 1917 483">Docker runs processes in isolated containers. A container is a process which runs on a host. The host may be local or remote. When you execute <code>docker run</code>, the container process that runs is isolated in that it has its own file system, its own networking, and its own isolated process tree separate from the host.</p> <p data-bbox="646 524 1230 557">https://docs.docker.com/engine/reference/run/</p>

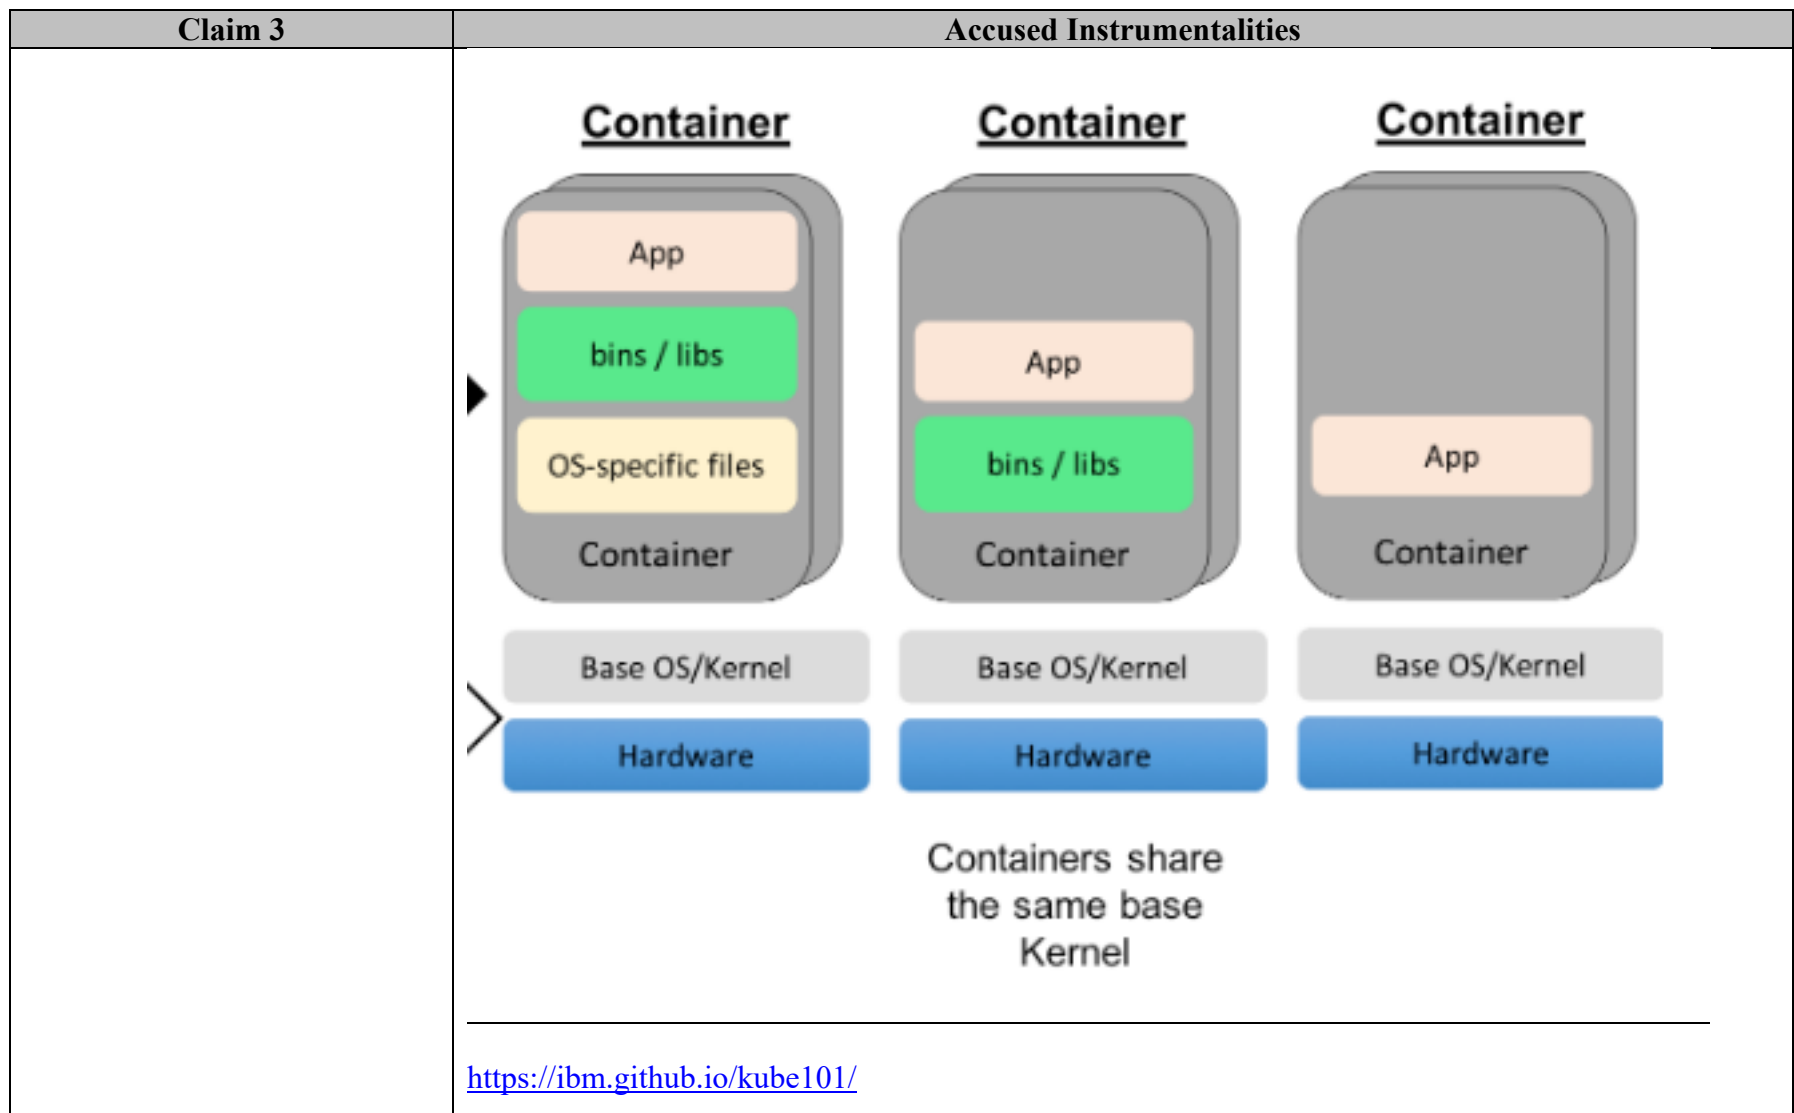
Claim 3

Claim 3	Accused Instrumentalities
<p data-bbox="205 735 611 946">3. A computing system according to claim 1 wherein OSCSEs corresponding to and capable of performing the same function as SLCSEs remain in the operating system kernel.</p>	<p data-bbox="646 735 1948 841">Each Accused Instrumentality comprises or constitutes a computing system according to claim 1 wherein OSCSEs corresponding to and capable of performing the same function as SLCSEs remain in the operating system kernel.</p> <p data-bbox="646 865 1927 930">For example, both Docker and Kubernetes systems preserve the host kernel substantially unchanged; therefore the OSCSEs corresponding to the SLCSEs remain in the operating system kernel.</p> <p data-bbox="646 963 762 995"><i>See, e.g.:</i></p> <p data-bbox="674 1044 1728 1068">A Docker image is the basis for every container that you create with IBM Cloud® Kubernetes Service.</p> <p data-bbox="674 1109 1902 1214">An image is created from a Dockerfile, which is a file that contains instructions to build the image. A Dockerfile might reference build artifacts in its instructions that are stored separately, such as an app, the app's configuration, and its dependencies.</p> <p data-bbox="646 1287 1461 1320">https://cloud.ibm.com/docs/containers?topic=containers-images</p>

Claim 3	Accused Instrumentalities		
	Docker base image	Supported versions	Source of security notices
	Alpine	All stable versions with vendor security support.	Alpine SecDB database  .
	Debian	All stable versions with vendor security support. CVEs on binary packages that are associated with the Debian source package <code>linux</code> , such as <code>linux-libc-dev</code> , are not reported. Most of these binary packages are kernel and kernel modules, which are not run in container images.	Debian Security Bug Tracker  .
	GoogleContainerTools distroless	All stable versions with vendor security support.	GoogleContainerTools distroless  .
	Red Hat® Enterprise Linux® (RHEL)	RHEL/UBI 7, RHEL/UBI 8, and RHEL/UBI 9	Red Hat Security Data API  .
	Ubuntu	All stable versions with vendor security support.	Ubuntu CVE Tracker  .
	https://cloud.ibm.com/docs/Registry?topic=Registry-va_index&interface=ui		

Claim 3	Accused Instrumentalities
	<p data-bbox="682 240 1045 289">Container images</p> <p data-bbox="682 316 1377 440">A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <hr data-bbox="644 459 1461 462"/> <p data-bbox="644 467 1241 500">https://kubernetes.io/docs/concepts/containers/</p> <p data-bbox="653 581 1551 764">Container image files are complete, static and executable versions of an application or service and differ from one technology to another. Docker images are made up of multiple layers, which start with a base image that includes all of the dependencies needed to execute code in a container. Each image has a readable/writable layer on top of static unchanging layers. Because each container has its own specific container layer that customizes that specific container, underlying image layers can be saved and reused in multiple containers. An Open Container Initiative (OCI)</p> <p data-bbox="644 776 1892 846">https://www.techtarget.com/searchitoperations/definition/container-containerization-or-container-based-virtualization</p>

Claim 3	Accused Instrumentalities
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p>https://docs.docker.com/storage/storagedriver/</p>



Claim 4

Claim 4	Accused Instrumentalities
<p>4. A computing system according to claim 1 wherein the one or more SLCSEs provided to one of the plurality of software applications having exclusive use thereof, use system calls to access services in the operating system kernel.</p>	<p>Each Accused Instrumentality comprises or constitutes a computing system according to claim 1 wherein the one or more SLCSEs provided to one of the plurality of software applications having exclusive use thereof, use system calls to access services in the operating system kernel.</p> <p>For example, the SLCSEs in a container use system calls to access services in the operating system kernel. For example, the glibc library (or other similar library) in the container uses system calls to interface with the host Linux kernel. In general, system calls can be observed using a tool such as strace.</p> <p><i>See, e.g.:</i></p> <p>The GNU C Library, commonly known as glibc, is the GNU Project implementation of the C standard library. It is a wrapper around the system calls of the Linux kernel for application use. Despite its name, it now also directly supports C++ (and, indirectly, other programming languages). It was started in the 1980s by the Free Software Foundation (FSF) for the GNU operating system.</p> <p>https://en.wikipedia.org/wiki/Glibc</p>

We can now get the process id directly from the cgroup. It will be located in the cgroup.procs file.

Terminal 2 - Worker Node

Get the process id

```
$ cat cri-containerd-ceeeef06afe89c8223d33b11e8d9e0b207118ac4dac3af826687668ee1ee16254
```

Validate what is running under the process

```
$ ps aux | grep 16254
```

```
azureus+ 16254 0.0 0.1 713972 10476 ? Ssl 15:04 0:00 ./faultyapp
azureus+ 94806 0.0 0.0 7004 2168 pts/0 S+ 16:22 0:00 grep --color=a
```

Got it! With that, we can try to find out what is going out inside the app. Lets try to run strace to get some more insight.

Terminal 2 - Worker Node

```
$ sudo strace -p 16254 -f
```

```
...
```

The app is trying to read a file port.txt

```
[pid 16269] openat(AT_FDCWD, "port.txt", O_RDONLY|O_CLOEXEC <unfinished ...>
```

```
[pid 16254] epoll_pwait(5, <unfinished ...>
```

The file does not exist

```
[pid 16269] <... openat resumed> = -1 ENOENT (No such file or directory)
```

```
[pid 16254] <... epoll_pwait resumed>[, 128, 0, NULL, 0) = 0
```

```
[pid 16269] write(1, "Something went wrong...\n", 24 <unfinished ...>
```

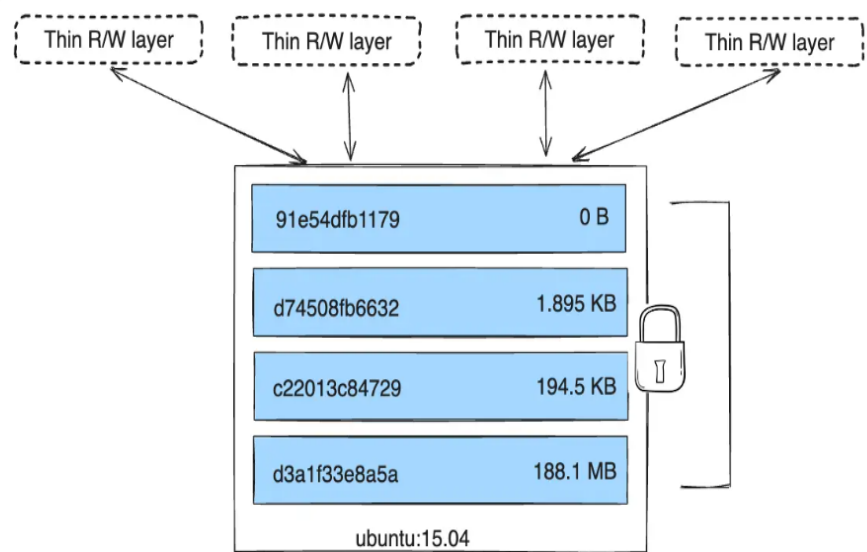
After filtering the output, we can see the application is trying to read a text file called port.txt, and a few lines later, there is a message stating ENOENT (No such file or directory). Let's create that file.

<https://www.berops.com/blog/a-different-method-to-debug-kubernetes-pods>

Claim 18

Claim 18	Accused Instrumentalities
<p>18. A computer system as defined in claim 2 wherein SLCSEs are not copies of OSCSEs.</p>	<p>Each Accused Instrumentality comprises or constitutes a computer system as defined in claim 2 wherein SLCSEs are not copies of OSCSEs.</p> <p>For example, in a typical case the SLCSEs come from a Linux distribution independent of the host operating system, and thus are not identical to the OSCSEs. For another example, the SLCSEs are provided to the computer system through a separate process than the process by which the OSCSEs are provided to the computer system, and thus are not copied from the OSCSEs.</p> <p><i>See, e.g.:</i></p> <p>Containers are executable units of software in which application code is packaged along with its libraries and dependencies, in common ways so that the code can be run anywhere—whether it be on desktop, traditional IT or the cloud.</p> <p>To do this, containers take advantage of a form of operating system (OS) virtualization in which features of the OS kernel (e.g. Linux namespaces and cgroups, Windows silos and job objects) can be leveraged to isolate processes and control the amount of CPU, memory and disk that those processes can access.</p> <p>Containers are small, fast and portable because unlike a virtual machine, containers do not need to include a guest OS in every instance and can instead simply leverage the features and resources of the host OS.</p> <p>https://www.ibm.com/topics/containers</p>

Claim 18	Accused Instrumentalities
	<p data-bbox="684 256 1875 505">Containers use a form of operating system (OS) virtualization. Put simply, they leverage features of the host operating system to isolate processes and control the processes' access to CPUs, memory and desk space.</p> <p data-bbox="644 570 1241 602">https://www.ibm.com/blog/containers-vs-vms/</p> <h2 data-bbox="684 643 1075 699">Container images</h2> <p data-bbox="684 732 1428 878">A container image is a ready-to-run software package containing everything needed to run an application: the code and any runtime it requires, application and system libraries, and default values for any essential settings.</p> <p data-bbox="644 911 1241 943">https://kubernetes.io/docs/concepts/containers/</p> <p data-bbox="657 995 1707 1157">Docker is used to create, run and deploy applications in containers. A Docker image contains application code, libraries, tools, dependencies and other files needed to make an application run. When a user runs an image, it can become one or many instances of a container.</p> <p data-bbox="644 1187 1547 1219">https://www.techtarget.com/searchitoperations/definition/Docker-image</p>

Claim 18	Accused Instrumentalities
	<p data-bbox="657 245 1911 375">Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p>  <p data-bbox="646 1019 1236 1052">https://docs.docker.com/storage/storagedriver/</p>